

TESSELLATIONS OF HYPERBOLIC SURFACES

JASON DEBLOIS

ABSTRACT. A finite subset S of a closed hyperbolic surface F canonically determines a *centered dual decomposition* of F : a cell structure with vertex set S , geodesic edges, and 2-cells that are unions of the corresponding Delaunay polygons. Unlike a Delaunay polygon, a centered dual 2-cell Q is not determined by its collection of edge lengths; but together with its combinatorics, these determine an *admissible space* parametrizing geometric possibilities for the Delaunay cells comprising Q . We illustrate its application by using the centered dual decomposition to extract combinatorial information about the Delaunay tessellation among certain genus-2 surfaces, and with this relate injectivity radius to covering radius here.

A finite subset \mathcal{S} of a closed hyperbolic surface F canonically determines a *Voronoi tessellation* V and *Delaunay tessellation* P , polygonal decompositions of F that are dual in a certain sense. Let us briefly outline this construction. Fix a locally isometric universal covering $\pi: \mathbb{H}^2 \rightarrow F$ and let $\tilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$. The *Voronoi tessellation* \tilde{V} of \mathbb{H}^2 determined by $\tilde{\mathcal{S}}$ is a cell complex structure where each $x \in \mathcal{S}$ determines a polygonal 2-cell V_x defined by:

$$(0.0.1) \quad V_x = \{p \in \mathbb{H}^2 \mid d(p, x) \leq d(p, y) \text{ for each } y \in \mathcal{S} - \{x\}\}$$

Then $\tilde{V}^{(1)} = \bigcup \{V_x \cap V_y \mid x \in \tilde{\mathcal{S}}, y \in \tilde{\mathcal{S}} - \{x\}\}$, and each point of $\tilde{V}^{(0)}$ is equidistant from at least 3 points of $\tilde{\mathcal{S}}$ (see Section 1). The *geometric dual* to an edge $V_x \cap V_y$ of \tilde{V} is the geodesic arc γ_{xy} in \mathbb{H}^2 joining x to y , and the set of geometric duals to edges of \tilde{V} is the edge set of the Delaunay tessellation \tilde{P} determined by $\tilde{\mathcal{S}}$. The covering action of $\pi_1 F$ on \mathbb{H}^2 leaves \tilde{V} and \tilde{P} invariant, and these descend to the tessellations V and P of F .

P and V are dual in the sense that their edge sets are canonically bijective, as is the vertex set of each with the face set of the other. However, in some cases an edge $e = V_x \cap V_y$ of \tilde{V} is not *centered* (see Definition 3.1): $\text{int } e$ does not intersect the geometric dual $\gamma_{xy} \subset \tilde{P}$ to e . We will regard this as a pathology of P , and “fix” it with the centered dual decomposition. Before we outline this construction, here is a sample application of our methods:

Theorem 0.1. *Let $r_\beta = d_\beta/2 > 0$, where $\cosh d_\beta$ is the real root of $x^3 - 14x^2 - 15x - 4$. The Delaunay tessellation of a closed, orientable hyperbolic surface F of genus 2 determined by $\{x\}$ has all edges centered if F has injectivity radius $r \geq r_\beta$ at x . It is a triangulation unless $r = r_\beta$ and each edge has length d_β , in which case it has a single quadrilateral 2-cell.*

The numerical value of r_β is roughly 1.7006, whereas Boröczky’s theorem [1] implies a universal upper bound of $r_\alpha \cong 1.7191$ on the injectivity radius of a genus-2 hyperbolic surface at a point x (see Lemma 2.3). Example 2.2 describes a surface F_α with injectivity radius r_α

Partially supported by NSF grant DMS-1007175.

at some $x_\alpha \in F_\alpha$, showing that Boröczky's upper bound is sharp. Example 2.4 describes a surface F_β with injectivity radius r_β at $x_\beta \in F_\beta$ and a quadrilateral 2-cell in the Delaunay tessellation determined by $\{x_\beta\}$, and Example 2.12 describes arbitrarily small deformations of F_β with Delaunay tessellations that have non-centered edges.

For surfaces satisfying its hypotheses, Theorem 0.1 has the following geometric consequence:

Theorem 0.2. *The space $\mathcal{M}_2^{(r_\beta)}$ of closed, orientable hyperbolic surfaces F with injectivity radius at least r_β at some $x \in F$ is compact. If $F \in \mathcal{M}_2^{(r_\beta)}$ has injectivity radius $r \geq r_\beta$ at x , the covering radius J of F at x satisfies $\sinh J \leq \sqrt{2} \sinh r$.*

We regard $\mathcal{M}_2^{(r_\beta)}$ above as a subspace of the moduli space \mathcal{M}_2 of closed, orientable hyperbolic surfaces of genus 2, given its usual topology, see eg. [4]. By the *covering radius* of F at x we refer to the infimum of $r > 0$ such that F is contained in the open r -neighborhood of x . The bound above is sharp, realized on F_β from Example 2.4, at x_β .

Let us briefly recall the well-known Mumford compactness criterion [5], that for any $r > 0$, the set of surfaces with injectivity radius at least r at *every* point is compact in \mathcal{M}_2 . The analogous generalization of Corollary 0.2 does not hold, since if F has small enough injectivity radius at x it can be deformed by making a distant curve arbitrarily short while keeping the injectivity radius at x constant. Note that this increases the covering radius at x to infinity.

The main “result” of the paper is really the construction of a *centered dual tessellation*, and the attachment of *admissible spaces* to its 2-cells, a process we now outline. In Section 1 we introduce terminology and define the Voronoi and Delaunay tessellations determined by a finite subset \mathcal{S} of a surface F . This material is standard. Section 2 gives a series of examples to motivate what follows, including a Delaunay tessellation that is not a triangulation (see Corollary 2.11), and another with a non-centered edge (see Lemma 2.13).

Section 3 gives deeper information on this failure of duality. If an edge e of the Voronoi tessellation V is not centered, one may orient it pointing “away” from γ . Lemma 3.3 asserts that a Delaunay 2-cell P_v is centered if and only if the associated vertex $v \in \mathcal{V}^{(0)}$ is not the initial vertex of any non-centered edge. Moreover each component of the union $\mathcal{V}_n^{(1)}$ of non-centered edges is a tree with a canonical root vertex, by Lemma 3.6.

We define the *centered dual graph* $P_c^{(1)}$ to the Voronoi tessellation to be the union of edges of P geometrically dual to centered edges of V , and show in the remainder of Section 3 that $P_c^{(1)}$ is the one-skeleton of a cell decomposition P_c , the *centered dual decomposition* of F , with vertex set \mathcal{S} (see Definition 3.18). By Proposition 3.16, each 2-cell Q of P_c is the union of Delaunay cells P_v such that $v \in Q \cap V^{(0)}$. This is either $T^{(0)}$ for a component T of $V_n^{(1)}$ or a vertex v not contained in any such component, by Lemma 3.12.

Each 2-cell of the Delaunay tessellation is *cyclic*: all its vertices are equidistant from the center of the corresponding vertex of V (see Lemma 1.2). It follows that such cells are each determined up to isometry by their side length collections [6] (cf. [7] or [3]). This does not hold for a centered dual 2-cell Q containing a component T of $V_n^{(1)}$. Instead, in Section 5 we will describe an *admissible space* $Ad(\mathbf{d}_\mathcal{F})$, determined by T and the side length collection

$\mathbf{d}_{\mathcal{F}}$ of Q , that in some sense parametrizes all possible combinations of Delaunay cells that can comprise Q with side length collection $\mathbf{d}_{\mathcal{F}}$. In particular see Lemma 5.4.

Our main application of the centered dual/admissible space construction is a machine for turning lower bounds on the side lengths of a centered dual 2-cell with few edges into a good lower bound on its area, described in Section 6. The corresponding problem for Delaunay 2-cells is complicated by non-centeredness: as we observed in [3], the area of a non-centered cyclic polygon decreases as the length of its longest side increases.

By Lemma 5.7, given a rooted tree T and side length collection $\mathbf{d}_{\mathcal{F}}$, the sum of areas of the Delaunay polygons comprising the corresponding centered dual 2-cell determines a continuous function on the closure $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$ of $Ad(\mathbf{d}_{\mathcal{F}})$. We show that if T has only one or two edges, the minimum of this function occurs at one of a few tightly-prescribed places. Section 6 describes an algorithm that produces lower bounds for the values at such locations, given lower bounds on the coordinates of $\mathbf{d}_{\mathcal{F}}$.

In fact we work in Sections 5 and 6 with the *radius- R defect*, which in best cases records the area of the region in a polygon but outside the union of disks of radius R centered at its vertices. Section 4 is devoted to establishing Proposition 4.2, which asserts that this does hold for a centered dual 2-cell Q . This proves convenient in applications.

In Section 7 we make some computations and prove Theorem 0.1. We prove Theorem 0.2 in Section 8.

Acknowledgments. This paper answers an objection of Marc Culler. The author thanks Peter Shalen and Steve Kerckhoff for helpful conversations.

1. THE VORONOI AND DELAUNAY TESSELLATIONS

This section gives a self-contained introduction to the Voronoi and Delaunay tessellations determined by a finite subset of a hyperbolic surface. Let us first establish some notation.

If γ is a geodesic in \mathbb{H}^2 , a *half-space bounded by γ* is the closure of a component of $\mathbb{H}^2 - \gamma$. We will say the *frontier* of $K \subset \mathbb{H}^2$ is $fr(K) \doteq K \cap \overline{\mathbb{H}^2 - K}$. Thus for instance γ is the frontier of either half-space bounded by γ . If $\{\mathcal{H}_i\}$ is a collection of half-spaces, each with bounding geodesic γ_i , we say that $P = \bigcap_i \mathcal{H}_i$ is a *convex polygon* if it is nonempty and the collection $\{\gamma_i\}$ is locally finite; ie, for each $p \in P$ there is an open set U and a finite collection $\{\gamma_{i_1}, \dots, \gamma_{i_n}\}$ of boundary geodesics such that $U \cap (\bigcup \gamma_i) \subset \bigcup_{j=1}^n \gamma_{i_j}$.

An *edge* (or *side*) of $P = \bigcap \mathcal{H}_i$ is $\gamma_i \cap P$ for some i such that this intersection is non-empty or a singleton, and the *boundary* ∂P of P is the union of its edges. One finds that ∂P is the topological frontier $P \cap \overline{\mathbb{H}^2 - P}$ of P in \mathbb{H}^2 . A *vertex* of P is the nonempty intersection of two edges. We will say P is *cyclic* if all its vertices are equidistant from some $v \in \mathbb{H}^2$, the *center* of P , and P is *centered* if $v \in \text{int } P$. The *radius* of a cyclic polygon P is J such that $d(v, x) = J$ for each vertex x of P , where v is the center of P .

By the *injectivity radius* of $\mathcal{S} \subset \mathbb{H}^2$ we refer to the supremum of the set of $r \geq 0$ such that $d(x, y) > 2r$ for all distinct x and y in \mathcal{S} . If \mathcal{S} has injectivity radius $R > 0$, then for any distinct x and y in \mathcal{S} the open disk $B_R(x) \doteq \{p \in \mathbb{H}^2 \mid d(x, p) < R\}$ is disjoint from $B_R(y)$.

Fact. If $\mathcal{S} \subset \mathbb{H}^2$ has injectivity radius $r > 0$ then $\mathcal{S} \cap K$ is finite for any bounded set $K \subset \mathbb{H}^2$.

This is because the R -neighborhood of K has finite area and so cannot contain infinitely many disjoint disks with a fixed positive area. It follows that if $\mathcal{S} \subset \mathbb{H}^2$ has positive injectivity radius then it is closed and discrete. The converse is not true, but a closed and discrete set \mathcal{S} does satisfy the fact above. This will suffice to define the Voronoi and Delaunay tessellations.

For distinct x and y in \mathbb{H}^2 , we will often refer by γ_{xy} to the unique geodesic arc joining x to y , and by γ_{xy}^\perp to its *perpendicular bisector*: the hyperbolic geodesic intersecting γ_{xy} at its midpoint m , at right angles. For each $p \in \gamma_{xy}^\perp$ the hyperbolic law of cosines gives:

$$\cosh d(x, p) = \cosh d(x, m) \cosh d(p, m) = \cosh d(y, p)$$

Thus the points of γ_{xy}^\perp are equidistant from x and y ; in fact $\gamma_{xy}^\perp = \{p \in \mathbb{H}^2 \mid d(p, x) = d(p, y)\}$.

Lemma 1.1. *If $\mathcal{S} \subset \mathbb{H}^2$ is closed and discrete then for each $x \in \mathcal{S}$, V_x as defined in (0.0.1) is a convex polygon in \mathbb{H}^2 , and if \mathcal{S} has injectivity radius $R > 0$ then $\overline{B_R(x)} \subset V_x$.*

Proof. For $y \in \mathcal{S} - \{x\}$, let \mathcal{H}_{xy} be the half-space containing x and bounded by γ_{xy}^\perp . Then $\mathcal{H}_{xy} = \{p \mid d(p, x) \leq d(p, y)\}$. It follows immediately that (0.0.1) may be rewritten as:

$$V_x = \bigcap_{y \in \mathcal{S} - \{x\}} \mathcal{H}_{xy}$$

Fix $p \in V_x$ and let $J = d(x, p)$. If $K = B_{3J}(p)$ then as we pointed out above the lemma, $K \cap \mathcal{S}$ is a finite set $\{x, y_1, \dots, y_n\}$. If $U = B_{J/2}(p)$, then for $q \in U$ the triangle inequality gives $d(x, q) < 3J/2$, so for y outside of $B_{3J}(p)$ it follows that $\gamma_{xy}^\perp \cap U = \emptyset$. Therefore:

$$U \cap \left(\bigcup_{y \in \mathcal{S} - \{x\}} \gamma_{xy}^\perp \right) \subset \bigcup_{i=1}^n \gamma_{xy_i}^\perp,$$

and V_x is a convex polygon.

If \mathcal{S} has injectivity radius $R > 0$, then since the open disk $B_R(x)$ does not intersect $B_R(y)$ for any $y \in \mathcal{S} - \{x\}$, the points of $B_R(x)$ are closer to x than any $y \in \mathcal{S} - \{x\}$. Thus $B_R(x) \subset V_x$; in particular, V_x is nonempty, and since it is closed it contains $\overline{B_R(x)}$. \square

It is easy to show that $\text{int } V_x = \{p \mid d(p, x) < d(p, y) \text{ for each } y \in \mathcal{S} - \{x\}\}$, and that $\partial V_x = \bigcup_{y \in \mathcal{S} - \{x\}} V_x \cap V_y$. Furthermore, for any $y \in \mathcal{S} - \{x\}$ such that $V_x \cap V_y$ is nonempty, it is contained in the equidistant locus γ_{xy}^\perp . This gives:

Fact. If $\mathcal{S} \subset \mathbb{H}^2$ is closed and discrete, then for distinct x, y , and z in \mathcal{S} , $V_x \cap V_y \cap V_z$ contains at most a single point.

This is because each point of $V_x \cap V_y \cap V_z$ is in both γ_{xy}^\perp and γ_{xz}^\perp , and since $y \neq z$ these are distinct geodesics which therefore meet transversely in a single point (if at all).

Lemma 1.2. *For $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete, define $V^{(0)} = \{V_x \cap V_y \cap V_z \mid x, y, z \in \mathcal{S} \text{ distinct}\}$. For $v \in V^{(0)}$, there exists $J_v > 0$ such that $v \in V_x$ if and only if $d(v, x) = J_v$ for each $x \in \mathcal{S}$. In particular, $\mathcal{S} \cap B_{J_v}(v) = \emptyset$.*

Proof. Suppose $x \in \mathcal{S}$ has $v \in V_x$, and let $J_v = d(x, v)$. It follows directly from the definition (0.0.1) that $d(v, y) \geq J_v$ for all $y \in \mathcal{S} - \{x\}$. For some such y , if $v \in V_y$ then $d(v, x) \geq d(v, y)$ by the definition of V_y , so we must have $d(v, y) = J_v$. This proves the lemma. \square

Corollary 1.3. *For $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete, $V^{(0)}$ from Lemma 1.2 is closed and discrete.*

Proof. For $v \in V^{(0)}$ let $J_v > 0$ be prescribed by Lemma 1.2. Each point of $B_{J_v}(v)$ is within $2J_v$ of a point of \mathcal{S} , so for $v' \in V^{(0)} \cap B_{J_v}(v)$ we have $J_{v'} < 2J_v$. Thus for such v' , $B_{3J_v}(v)$ contains the set of $x \in \mathcal{S}$ such that $v' \in V_x$. The fact above Lemma 1.2 implies that v' is determined by this set, so since $\mathcal{S} \cap B_{3J_v}(v)$ is finite there can be only finitely many $v' \in V^{(0)} \cap B_{J_v}(v)$. \square

Fact. For $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete, $x \in \mathcal{S}$, and each edge $e = V_x \cap V_y$ of V_x , $\text{int}(e) \cap V_z = \emptyset$ for each $z \in \tilde{\mathcal{S}} - \{x, y\}$.

This is because $v = e \cap V_z$ is an intersection of the geodesics γ_{xy}^\perp and γ_{xz}^\perp , being equidistant from x, y , and z , and so one interval of $\gamma_{xy}^\perp - v$ consists entirely of points closer to z than x .

By the fact above, the set of edges of Voronoi polygons has the structure of an embedded graph in \mathbb{H}^2 with vertex set $V^{(0)}$.

Definition 1.4. For $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete, the *Voronoi tessellation* V determined by \mathcal{S} is the cell complex structure with 2-cells of the form V_x for $x \in \mathcal{S}$, and with $V^{(1)} = \bigcup \{V_x \cap V_y \mid \text{distinct } x, y \in \mathcal{S}\}$ and $V^{(0)} = \bigcup \{V_x \cap V_y \cap V_z \mid x, y, z \in \mathcal{S} \text{ distinct}\}$.

We caution that this definition does not imply that the Voronoi tessellation has trivalent one-skeleton; only that each vertex is contained in at least three 2-cells.

Definition 1.5. Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete. For $v \in V^{(0)}$, say the collection $\{e_0, \dots, e_{n-1}\}$ of edges of V containing v is *cyclically ordered* if for each i there exists $x_i \in \mathcal{S}$ so that e_i and e_{i+1} are edges of V_{x_i} (taking $i+1$ modulo n).

Note that if the edges containing $v \in V^{(0)}$ are cyclically ordered e_0, \dots, e_{n-1} then $e_i = V_{x_i} \cap V_{x_{i-1}}$ for each i (with $i-1$ taken modulo n), where the x_i are as in the definition above.

Lemma 1.6. *Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete, and fix $v \in V^{(0)}$. Let $C = \{p \in \mathbb{H}^2 \mid d(v, p) = J_v\}$, let the edges containing v be cyclically ordered $\{e_0, \dots, e_{n-1}\}$ as in Definition 1.5, and let $\{x_i\} \subset \mathcal{S}$ be the associated collection with e_i and $e_{i+1} \subset V_{x_i}$ for each i . Then $\{x_i\}_{i=0}^{n-1} = \{x \in \mathcal{S} \mid v \in V_x\}$ is cyclically ordered in the sense of [3, Definition 1.3].*

Proof. For $y \in \mathcal{S}$, if $d(v, y) = J_v$ then v is as close to y as to any other element of \mathcal{S} ; hence $v \in V_y$. Since $d(v, y) = J_v = d(v, x_i)$ for each i , v is in the frontier of V_y and hence is in an




FIGURE 1.1. Some objects from the proof of Lemma 1.6. (Edges of V are bold.)

edge e of V_y . By hypothesis there is an i such that $e = e_i$, so by the observation above the lemma $y = x_i$ or x_{i-1} . This shows that $\{x \in \mathcal{S} \mid v \in V_x\} = \{x_i\}_{i=0}^{n-1}$.

For each i , since V_{x_i} is convex it contains the geodesic arc λ_i joining x_i to v . Now fix some $i \in \{0, \dots, n-1\}$ and a point $w \in e_i \cap B_{J_v}(v)$ near v . Then by convexity of V_{x_i} , the triangle T_i determined by v , w , and x_i is entirely contained in V_{x_i} . Also $T_i \subset \overline{B_{J_v}(v)}$, again by convexity (since $x_i \in C = \partial B_{J_v}(v)$). By the same argument, the triangle T_{i-1} determined by v , w , and x_{i-1} is contained in $V_{x_{i-1}} \cap \overline{B_{J_v}(v)}$. (See Figure 1.1.)

Let α_{i-1} be the vertex angle of T_{i-1} at v ; since $w \in \text{int}(e_i)$ this is the angle in $V_{x_{i-1}}$ between λ_{i-1} and e_i . Similarly, the angle α_i of T_i at v is the angle in V_{x_i} between e_i and λ_i . If $\epsilon > 0$ is less than the distance from v to the side $[x_i, w]$ of T_i joining x_i and w , then T_i contains the entire sector of $B_\epsilon(v)$, with angle α_i , determined by e_i and λ_i . If $\epsilon < d(v, [x_{i-1}, w])$ then the analogous assertion holds for the sector of $B_\epsilon(v)$ determined by λ_{i-1} and e_i .

Fix $\epsilon > 0$ satisfying the requirements of the paragraph above and let C_ϵ be the circle of radius ϵ centered at v . The above implies that $C_\epsilon \cap (T_{i-1} \cup T_i)$ is an interval of C_ϵ with angle measure $\alpha_{i-1} + \alpha_i$ and endpoints $\lambda_{i-1} \cap C_\epsilon$ and $\lambda_i \cap C_\epsilon$. There is a homeomorphism $C \rightarrow C_\epsilon$ that takes $x \in C$ to $C_\epsilon \cap \lambda_x$, where λ_x is the geodesic arc joining x to v . In particular, x_j maps to $\lambda_j \cap C_\epsilon$ for each $j \in \{0, \dots, n-1\}$. Let I be the preimage of $C_\epsilon \cap (T_{i-1} \cup T_i)$ in C , a closed subinterval bounded by x_{i-1} and x_i . For $j \neq i$ or $i-1$, x_j maps to a point outside $C_\epsilon \cap (T_{i-1} \cup T_i)$ since $\lambda_j \subset V_{x_j}$, and so $x_j \in C - I$, a component of $C - \{x_{i-1}, x_i\}$. \square

We will define the Delaunay tessellation P determined by $\mathcal{S} \subset \mathbb{H}^2$ as a sort of “dual” to the Voronoi tessellation determined by \mathcal{S} . In particular we take $P^{(0)} = \mathcal{S}$, in 1–1 correspondence with the set of 2-cells of V . The edges of P are also determined by edges of V :

Definition 1.7. Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete. For an edge $e = V_x \cap V_y$ of V , the *geometric dual* to e is the geodesic arc γ_{xy} joining x to y .

We define $P^{(1)}$ to be the union of geometric duals to edges of V . The lemma below establishes that $P^{(1)}$ has the structure of an embedded graph in \mathbb{H}^2 with vertex set \mathcal{S} .

Lemma 1.8. Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete. If $e = V_x \cap V_y$ and $e' = V_{x'} \cap V_{y'}$ are distinct edges of V , then their geometric duals satisfy $\gamma_{xy} \cap \gamma_{x'y'} = \{x, y\} \cap \{x', y'\}$.

Proof. If $x = x'$, say, then $y \neq y'$ since e and e' are distinct, so since γ_{xy} and $\gamma_{x'y'}$ are geodesic arcs they intersect only at x . We will thus assume that $\{x, y\} \cap \{x', y'\} = \emptyset$.

If e and e' share a vertex v , then upon cyclically ordering the edges containing v as e_0, \dots, e_{n-1} as in Definition 1.5, we have $e = e_i$ and $e' = e_{i'}$ for distinct i and i' in $\{0, \dots, n-1\}$. Then $\{x, y\} = \{x_{i-1}, x_i\}$ and $\{x', y'\} = \{x_{i'-1}, x_{i'}\}$, as we observed above Lemma 1.6, and Lemma 1.2 implies that γ_{xy} and $\gamma_{x'y'}$ are chords of the circle C of radius J_v centered at v .

Chords of C with distinct endpoints intersect if and only if the endpoints of one separate the endpoints of the other on C . But Lemma 1.6 implies that $x_{i'-1}$ and $x_{i'}$ share a component of $C - \{x_i, x_{i-1}\}$, so $\gamma_{xy} \cap \gamma_{x'y'} = \emptyset$ in this case. We therefore assume below that $e \cap e' = \emptyset$.

Let v be the nearer of the two endpoints of e to x and y , and let D_v be the disk of radius J_v centered at v . The same construction yields v' and a disk $D_{v'}$ of radius $J_{v'}$ associated to $\gamma_{x'y'}$. By Lemma 1.2, γ_{xy} is a chord of the circle ∂D_v , and $\gamma_{x'y'}$ is a chord of $\partial D_{v'}$. If γ_{xy} intersects $\gamma_{x'y'}$, their intersection point is contained in $D_v \cap D_{v'}$.

Claim 1.8.1. If distinct circles in \mathbb{H}^2 have intersecting chords with distinct endpoints, then one chord has an endpoint in the open disk in \mathbb{H}^2 complementary to the other circle.

Proof of claim. Let C and C' be distinct circles with chords γ and γ' , respectively, that intersect. If C is contained in the open disk determined by C' , then the claim is immediate. The same holds if C' is contained in the open disk determined by C , so we will assume that neither of these possibilities occurs. Then C must intersect C' , since each chord of C is contained in the disk that it bounds, and similarly for C' .

If $C \cap C'$ is a single point, this must also be $\gamma \cap \gamma'$, an endpoint of each, so this cannot occur. It follows that $C \cap C'$ consists of two points. We may assume that neither endpoint of γ is contained in the open disk complementary to C' , since otherwise the claim holds. Since γ intersects γ' , it nonetheless intersects C' . If $\gamma \cap C'$ is a single point, then this is also $\gamma \cap \gamma'$, an endpoint of γ' and therefore not an endpoint of γ . Since this endpoint of γ' is in the interior of γ , the claim holds in this case.

Let us now suppose that $\gamma \cap C'$ consists of two points, and let γ_0 be the closed subarc of γ bounded by $\gamma \cap C'$. Then $\gamma \cap \gamma' \subset \gamma_0$. The geodesics γ and γ' intersect transversely, so γ' has one endpoint in each of the open sub-arcs of C' complementary to $C' \cap \gamma$. Since one of these is contained in the open disk complementary to C , the claim holds. \square

Using the claim, we may assume that the endpoint x of γ_{xy} is contained in the open disk with radius $J_{v'}$ centered at v' . But this contradicts Lemma 1.2, and the result follows. \square

The 2-cells of the Delaunay tessellation are associated to $V^{(0)}$ by the lemma below. Recall that a compact, convex polygon P is *cyclic* if its vertices are equidistant from a fixed point, its *center* (see [3]), and that the *radius* of P is the distance from its center to the vertices.

Lemma 1.9. *Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete. For each $v \in V^{(0)}$ there is a cyclic polygon P_v in \mathbb{H}^2 with center v and radius J_v (as supplied by Lemma 1.2), such that:*

- If the edges of V containing v are cyclically ordered e_0, \dots, e_{n-1} , the vertex set of P_v is the collection $\{x_i\}_{i=0}^{n-1}$ from Definition 1.5.
- The edge set of P_v , cyclically ordered in the sense of [3, Definition 2.5], is $\{\gamma_i\}_{i=0}^{n-1}$, where γ_i is the geometric dual to e_i for each i . Furthermore, $P_v \cap P^{(1)} = \gamma_0 \cup \dots \cup \gamma_{n-1}$.

For $v \neq w$, $\text{int } P_v \cap P_w = \emptyset$, and P_v and P_w share an edge if and only if v and w are opposite endpoints of an edge of V .

Proof. Let v and the collections $\{e_0, \dots, e_{n-1}\}$, and $\{\gamma_0, \dots, \gamma_{n-1}\}$ be as described in the hypotheses of the lemma. The observation above Lemma 1.6 implies that for each i , γ_i joins x_{i-1} to x_i , where $\{x_i\}_{i=0}^{n-1} \subset \mathcal{S}$ is the collection from Definition 1.5. Since the collection $\{x_i\}$ is cyclically ordered by Lemma 1.6, Lemma 1.4 of [3] asserts there is a cyclic n -gon P_v center v , radius J_v , vertex set $\{x_i\}$ and edge set $\{\gamma_i\}$. Furthermore, since the x_i are cyclically ordered, the γ_i are as well (see [3, Definition 2.5]).

For any x and $y \in \mathcal{S}$, since V_x and V_y are convex their intersection is connected. This implies in particular that for i and $j \in \{0, \dots, n-1\}$, if V_{x_i} shares an edge e with V_{x_j} then $v \in e$. Since the only edges of V_{x_i} that contain v are e_i and e_{i+1} , it follows that $V_{x_i} \cap V_{x_j} = \{v\}$ unless $j = i \pm 1$ or i . Therefore by definition, no edge of $P^{(1)}$ joins x_i to x_j for $j \neq i \pm 1 \pmod{n}$. It thus follows from Lemma 1.8 that $P_v \cap P^{(1)} = \bigcup_{i=0}^{n-1} \gamma_i$.

If $p \in \text{int } P_v \cap P_w$, then since $\text{int } P_v \cap P^{(1)} = \emptyset$ and $\partial P_w \subset P^{(1)}$, $p \in \text{int } P_w$ and each geodesic ray from p intersects ∂P_v at or nearer to p than its point of intersection with ∂P_w . But since $\text{int } P_w \cap P^{(1)} = \emptyset$ and $\partial P_v \subset P^{(1)}$, each such point is in ∂P_w . It follows that $\partial P_v = \partial P_w$, and hence that $P_v = P_w$ (again see [3, Lemma 1.4]). This implies that $v = w$, since P_v is cyclic and a circle (and hence also its center) is determined by three points on it.

If P_v shares the edge γ of $P^{(1)}$ with P_w , then by construction the edge e of V dual to γ contains v and w . On the other hand, if v and w are vertices of an edge e of V , then again by construction the edge of $P^{(1)}$ dual to e is contained in P_v and P_w . \square

Definition 1.10. Suppose $\mathcal{S} \subset \mathbb{H}^2$ is closed and discrete. We take the *Delaunay tessellation* determined by \mathcal{S} to be the 2-complex P with vertex set \mathcal{S} , edge set the geometric duals to edges of the Voronoi tessellation V , and 2-cells P_v supplied by Lemma 1.9, for $v \in V^{(0)}$. For such v we will refer to P_v as the associated *vertex polygon*.

The Delaunay tessellation P is “dual” to the Voronoi tessellation V in the sense that there is a canonical one-to-one correspondence between its k -cells and the $(2-k)$ -cells of V for each $k \in \{0, 1, 2\}$. However, it is not necessarily dual in the sense of the intersection pairing: there is no reason in general that an edge of V should intersect its geometric dual, or that $v \in V^{(0)}$ should be in P_v .

The Delaunay tessellation is sometimes defined using “circumscribed circles,” but this has its problems in the hyperbolic setting, as the example below will demonstrate.




FIGURE 1.2. Combinatorial possibilities for the Voronoi and Delaunay tessellations determined by $\{x, y, z\} \subset \mathbb{H}^2$. (Edges of V are bold and of P , dashed.)

Example 1.11. Let $x, y, z \in \mathbb{H}^2$ be distinct, and suppose that $d(y, z) \geq \max\{d(x, y), d(x, z)\}$. It follows from [3, Lemma 2.4] that x, y , and z lie on a circle in \mathbb{H}^2 if and only if

$$(1.11.1) \quad \sinh(d(y, z)/2) < \sinh(d(x, y)/2) + \sinh(d(x, z)/2)$$

We motivate this fact with Figure 1.2, which uses the “upper half-plane” model for \mathbb{H}^2 : the set of complex numbers with positive imaginary coordinate, equipped with the hyperbolic Riemannian metric. It is well-known that in this model, each hyperbolic circle is also a Euclidean circle in \mathbb{C} , although with a different center and radius. However, some x, y , and z in \mathbb{H}^2 determine a Euclidean circle that does not lie entirely in \mathbb{H}^2 , as illustrated on the right-hand side of the figure, even if they do not lie on a hyperbolic geodesic.

If (1.11.1) does not hold, then the equidistant locus $V_x \cap V_y$ does not intersect $V_z \cap V_y$, as on the right-hand side of Figure 1.2. In this case V_x and V_z are each a single half-plane, V_y is bounded by two disjoint geodesics, and the Delaunay “tessellation” (as we have defined it) is the union of the dotted geodesic arcs γ_{xy} and γ_{xz} . If (1.11.1) does hold, then the Euclidean circle C containing x, y , and z lies in \mathbb{H}^2 and V_x, V_y , and V_z intersect at its (hyperbolic) center v . This is pictured on the left side of Figure 1.2, with the vertex polygon P_v shaded.

As in Example 1.11, the Delaunay tessellation determined by \mathcal{S} does not necessarily cover \mathbb{H}^2 ; indeed, if \mathcal{S} is finite then Lemma 1.9 implies that it is compact. However, we are primarily concerned here with tessellations that arise from closed hyperbolic surfaces — those which admit a locally isometric covering from \mathbb{H}^2 .

Lemma 1.12. *Let F be a closed hyperbolic surface, $\mathcal{S} \subset F$ a finite set, and $\pi: \mathbb{H}^2 \rightarrow F$ a locally isometric universal covering. Then $\tilde{\mathcal{S}} \doteq \pi^{-1}(\mathcal{S})$ has positive injectivity radius, and the Voronoi tessellation \tilde{V} and Delaunay tessellation \tilde{P} determined by $\tilde{\mathcal{S}}$ are invariant under the $\pi_1 F$ -action on \mathbb{H}^2 by covering transformations. Furthermore, V_x as defined in (0.0.1) is a compact polygon for each $x \in \tilde{\mathcal{S}}$, and \tilde{P} covers \mathbb{H}^2 .*

Proof. Since \mathcal{S} is finite and F is compact there is a lower bound $r > 0$ on the lengths of non-constant geodesic arcs in F with endpoints in \mathcal{S} . The injectivity radius of $\tilde{\mathcal{S}}$ is then $r/2 > 0$. Since $\tilde{\mathcal{S}}$ is invariant under the action of $\pi_1 F$, and this action is by isometries, it follows from (0.0.1) that $p \in V_x$ if and only if $g.p \in V_{g.x}$ for $p \in \mathbb{H}^2$, $x \in \tilde{\mathcal{S}}$, and $g \in \pi_1 F$. Therefore $g.V_x = V_{g.x}$, and $g.(V_x \cap V_y) = V_{g.x} \cap V_{g.y}$ for $y \in \tilde{\mathcal{S}} - \{x\}$. It follows that \tilde{V} is $\pi_1 F$ -invariant, and from this that \tilde{P} is as well.

Since F is compact, there exists $R > 0$ such that $\bigcup_{y \in \mathcal{S}} B(y, R)$ covers F (here $B(y, R)$ is the open R -neighborhood of y in the hyperbolic metric on F). Then $V_x \subset B(x, R) \subset \mathbb{H}^2$ for each $x \in \tilde{\mathcal{S}}$, and hence it is compact. In particular, V_x has only finitely many edges and vertices.

We will show that \tilde{P} is open and closed in \mathbb{H}^2 , and hence that it is all of \mathbb{H}^2 . First we claim that the collection of vertex polygons is locally finite: for any given v , Lemma 1.9 implies that $\text{int}(P_v)$ is disjoint from any other vertex polygon, and that the interior of an edge intersects exactly one other vertex polygon. Each vertex of P_v is some $x \in \tilde{\mathcal{S}}$, and since V_x has only finitely many vertices x is in only finitely many P_v . The claim follows, and therefore $\tilde{P} = \bigcup_{v \in \tilde{V}^{(0)}} P_v$ is closed in \mathbb{H}^2 .

We claim that \tilde{P} contains an open neighborhood of each $x \in \tilde{\mathcal{S}}$. For such x , enumerate the edges of V_x as e_0, \dots, e_{n-1} so that e_i intersects e_{i+1} in a vertex v_i for each i , taking $i+1$ modulo n . Then $x \in P_{v_i}$ for each i , and P_{v_i} intersects $P_{v_{i-1}}$ along the geometric dual γ_i to e_i and $P_{v_{i+1}}$ along the geometric dual γ_{i+1} to e_{i+1} (here again take $i \pm 1$ modulo n). For each i there exists $\epsilon_i > 0$ so that P_{v_i} intersects $\overline{B_x(\epsilon_i)}$ in the full sector determined by γ_i and γ_{i+1} , and we define $\epsilon = \min_i \{\epsilon_i\}$ and $C = \partial B_x(\epsilon)$. The claim follows from the fact that $C = \bigcup_{i=0}^{n-1} (P_{v_i} \cap C)$, since by the above this set is itself open and closed in C .

Edges of \tilde{V} are compact, so for each $v \in \tilde{V}^{(0)}$ and edge γ of P_v , Lemma 1.9 gives $\gamma = P_v \cap P_w$, where w is the other endpoint of the edge of \tilde{V} geometrically dual to γ . Therefore any point in the interior of γ has an open neighborhood in \mathbb{H}^2 contained in $P_v \cup P_w$. The claim above implies that each vertex of P_v also has an open neighborhood contained in \tilde{P} , and it follows that \tilde{P} is open in \mathbb{H}^2 . \square

Definition 1.13. For a closed hyperbolic surface F , a locally isometric universal covering $\pi: \mathbb{H}^2 \rightarrow F$, and $\mathcal{S} \subset F$ finite, let $\tilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$ and take $V = \pi(\tilde{V})$ and $P = \pi(\tilde{P})$ to be the *Voronoi tessellation* and *Delaunay tessellation determined by \mathcal{S}* , respectively, where \tilde{V} and \tilde{P} are as in Lemma 1.12.

Since a convex polygon is homeomorphic to a disk, and π takes the interior of each edge or 2-cell of \tilde{V} or \tilde{P} isometrically to F , V and P have the structure of cell decompositions of F . Note also that $\mathcal{S} = \pi(\tilde{\mathcal{S}})$ is the vertex set of P .

2. EXAMPLES AND TOOLS FOR RECOGNITION

In this section we will take advantage of tools from [3] for understanding the geometry of cyclic polygons, so let us begin by recalling some of its notation.

Definition 2.1 ([3], Definition 2.1). For $n \geq 3$, let $\sigma: \mathbb{R}^n \rightarrow \mathbb{R}^n$ be given by $\sigma(d_0, \dots, d_{n-1}) = (d_1, \dots, d_{n-1}, d_0)$, and refer by $\mathbb{R}^n / \mathbb{Z}_n$ to the quotient by the action of $\mathbb{Z}_n \doteq \langle \sigma \rangle$, and by

$[d_0, \dots, d_{n-1}]$ to the equivalence class in $\mathbb{R}^n/\mathbb{Z}_n$ of (d_0, \dots, d_{n-1}) . Define:

$$\begin{aligned}\widetilde{\mathcal{AC}}_n &= \left\{ (d_0, \dots, d_{n-1}) \in (\mathbb{R}^+)^n \mid \sinh(d_i/2) < \sum_{j \neq i} \sinh(d_j/2) \text{ for each } i \in \{0, \dots, n-1\} \right\} \\ \widetilde{\mathcal{C}}_n &= \left\{ (d_0, \dots, d_{n-1}) \in (\mathbb{R}^+)^n \mid \sum_{j=0}^{n-1} A_{d_j}(d_i/2) > 2\pi, \text{ where } d_i \geq d_j \forall j \in \{0, \dots, n-1\} \right\}\end{aligned}$$

Let $\mathcal{AC}_n = \widetilde{\mathcal{AC}}_n/\mathbb{Z}_n \subset \mathbb{R}^n/\mathbb{Z}_n$ and $\mathcal{C}_n = \widetilde{\mathcal{C}}_n/\mathbb{Z}_n \subset \mathcal{AC}_n$.

The point of this definition is that by [3, Proposition 2.7], each $(d_0, \dots, d_{n-1}) \in \widetilde{\mathcal{AC}}_n$ determines a cyclic n -gon with cyclically ordered side length collection given by its entries; this n -gon is unique up to isometry of \mathbb{H}^2 ; and two such points determine the same (oriented) n -gon if and only if they have the same class in \mathcal{AC}_n . We will say a cyclic n -gon is *represented* by $(d_0, \dots, d_{n-1}) \in \widetilde{\mathcal{AC}}_n$ if this tuple describes its cyclically ordered side length collection. [3, Proposition 2.7] further implies that each cyclic n -gon is represented by a point of $\widetilde{\mathcal{AC}}_n$.

The function $A_d(J)$ used in the definition of $\widetilde{\mathcal{C}}_n$ is defined in [3, Lemma 1.7]. Each point in $\widetilde{\mathcal{C}}_n$ determines a *centered* n -gon, a cyclic polygon P with center $v \in \text{int } P$ (recall from the beginning of Section 1 that the *center* of P is the center of the circle containing its vertices). Conversely, if $(d_0, \dots, d_{n-1}) \in \widetilde{\mathcal{AC}}_n$ represents a centered n -gon then it is in $\widetilde{\mathcal{C}}_n$. (It is not immediately obvious that $\widetilde{\mathcal{C}}_n \subset \widetilde{\mathcal{AC}}_n$, but this is proved in [3, Lemma 2.3].)

Example 2.2. Taking $D_0: \widetilde{\mathcal{AC}}_3 \rightarrow \mathbb{R}^+$ as in [3, Definition 5.1], determine $d_\alpha > 0$ by:

$$4\pi = 6 \cdot D_0(d_\alpha, d_\alpha, d_\alpha) = 6 \cdot D_{0,3}(d_\alpha) = 6 \cdot \left[\pi - 6 \sin^{-1} \left(\frac{1}{2 \cosh(d_\alpha/2)} \right) \right]$$

The latter equalities above follow from [3, Lemma 6.6]. Let $r_\alpha = d_\alpha/2$. Rearranging the equation above and taking sines of both sides gives:

$$\cosh r_\alpha = \frac{1}{2 \sin(\pi/18)} \cong 2.8794 \quad \cosh d_\alpha = \frac{1}{1 - \cos(\pi/9)} - 1 \cong 15.5817$$

By [3, Proposition 2.7] there is a centered triangle in \mathbb{H}^2 , unique up to isometry, with all side lengths d_α . Six copies T_1, \dots, T_6 of this triangle may be arranged in \mathbb{H}^2 so that they share a vertex and have disjoint interiors, and T_i shares an edge with $T_{i \pm 1}$ for $1 < i < 6$. Their union is thus an octahedron O_α , with all side lengths d_α and area 4π by construction. The Gauss-Bonnet formula implies that O_α has total angle defect 2π , so its quotient by some scheme for pairing edges that reverses boundary orientations and identifies all vertices is a genus-2 surface F_α . Let $x_\alpha \in F_\alpha$ be the projection of the vertices of O_α .

Since the angle measures of the T_i total 2π , each has angle $\pi/9$ at each vertex. For each i , open disks of radius r_α centered at the vertices of T_i do not intersect (see [3, Lemma 5.3]), and each intersects T_i in a full sector of angle measure $\pi/9$. Since the six non-overlapping T_i comprise O_α , a collection of open disks of radius r_α centered at each vertex of O_α intersects it in the non-overlapping union of 18 sectors of angle measure $\pi/9$. This projects to a

hyperbolic disk embedded in F_α , with center x_α and radius r_α . Since O_α has edge lengths $d_\alpha = 2r_\alpha$, F_α has injectivity radius r_α at x_α .

Boröczky's Theorem [1] implies that r_α is the largest injectivity radius possible at any point in any genus-two hyperbolic surface.

Lemma 2.3. *A closed, orientable hyperbolic surface F of genus 2 has injectivity radius at most r_α at any $x \in F$.*

Proof. Fix a locally isometric universal covering map $p: \mathbb{H}^2 \rightarrow F$. If F has injectivity radius R at x , then by definition it contains an isometrically embedded open hyperbolic disk D , with radius R , centered at x . Each point of $p^{-1}(x)$ is contained in a lift of D to \mathbb{H}^2 , and since D is embedded in F two such lifts do not overlap unless they are identical. Thus $p^{-1}(D)$ is a packing of \mathbb{H}^2 .

Let \tilde{V} be the Voronoi decomposition of \mathbb{H}^2 determined by $p^{-1}(x)$. Since $p^{-1}(x)$ has injectivity radius α , for any $\tilde{x} \in p^{-1}(x)$, $V_{\tilde{x}}$ contains the lift \tilde{D} of D centered at \tilde{x} (see Lemma 1.1). The main theorem of [1] implies:

$$\frac{\text{area}(\tilde{D})}{\text{area}(V_{\tilde{x}})} \leq d(R) = \frac{3\alpha(R) \cdot (\cosh R - 1)}{\pi - 3\alpha(R)}$$

Here $\alpha(R)$ is the vertex angle of an equilateral triangle in \mathbb{H}^2 with sides of length $2R$. Since $V_{\tilde{x}}$ projects onto F , isometrically on its interior, and \tilde{D} projects isometrically to D we have:

$$\frac{\text{area}(D)}{\text{area}(F)} \leq \frac{3\alpha(R) \cdot (\cosh R - 1)}{\pi - 3\alpha(R)}$$

Since F has area 4π and D has area $2\pi(\cosh R - 1)$, the above inequality simplifies to $\alpha(R) \geq \pi/9$. The hyperbolic law of cosines implies:

$$\cos \alpha(R) = \frac{\cosh^2(2R) - \cosh(2R)}{\sinh^2(2R)} = \frac{\cosh(2R)}{\cosh(2R) + 1} = 1 - \frac{1}{\cosh(2R) + 1}$$

Solving for $\cosh(2R)$ and applying the “half-angle” identities for the sine and hyperbolic cosine functions gives $\cosh R = 1/2 \sin(\alpha(R)/2) \leq 1/2 \sin(\pi/18)$. The conclusion follows. \square

Example 2.4. Let $d_\beta > 0$ be determined by the following criterion:

$$\begin{aligned} 4\pi &= 4 \cdot D_0(d_\beta, d_\beta, d_\beta) + D_0(d_\beta, d_\beta, d_\beta, d_\beta) = 4 \cdot D_{0,3}(d_\beta) + D_{0,4}(d_\beta) \\ &= 4 \cdot \left[\pi - 6 \sin^{-1} \left(\frac{1}{2 \cosh(d_\beta/2)} \right) \right] + \left[2\pi - 8 \sin^{-1} \left(\frac{\sqrt{2}}{2 \cosh(d_\beta/2)} \right) \right] \end{aligned}$$

The latter equalities follow from [3, Lemma 6.6]. Applying the identity $2 \sin^{-1} x = \cos^{-1}(1 - 2x^2)$ and the half-angle identity for hyperbolic cosine, and re-arranging yields:

$$\frac{\pi}{2} = 3 \cos^{-1} \left(\frac{\cosh d_\beta}{\cosh d_\beta + 1} \right) + \cos^{-1} \left(\frac{\cosh d_\beta - 1}{\cosh d_\beta + 1} \right)$$

After taking cosines of both sides and simplifying with trigonometric identities we find that $y = \cosh d_\beta + 1$ satisfies $y^3 - 17y^2 + 16y - 4$, and hence that $x = y - 1$ is as described in Theorem 0.1. With $r_\beta = d_\beta/2$ we have:

$$\cosh d_\beta \cong 15.0166 \qquad \cosh r_\beta \cong 2.8299$$

Let Q be a centered quadrilateral and T_1, T_2, T_3, T_4 centered triangles, each with all side lengths d_β . These exist by [3, Proposition 2.7]. Arrange them in \mathbb{H}^2 so that they are pairwise non-overlapping and T_i shares an edge with Q for each i . Then $O_\beta \doteq Q \cup (\bigcup_i T_i)$ is a hyperbolic octahedron with area 4π , and hence total angle defect 2π . An edge-pairing scheme as in Example 2.2 yields a genus-2 surface F_β , and arguing as in Example 2.2 we find that F_β has injectivity radius r_β at the point $x_\beta \in F_\beta$ descended from the vertices of O_β .

We now prove a few preliminary results that will allow us to pin down the Voronoi and Delaunay tessellations in Examples 2.2 and 2.4.

Lemma 2.5. *Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ with injectivity radius $R > 0$, and let $B_0 = \cosh^{-1}(2 \cosh(2R) - 1)$. For $x, y \in \mathcal{S}$, if $d(x, y) \leq B_0$ then the midpoint m of the geodesic arc γ_{xy} joining x to y is in $V_x \cap V_y$. If $d(x, y) < B_0$ then γ_{xy} is the geometric dual to an edge $e = V_x \cap V_y$ of V with $m = \gamma_{xy} \cap e \in \text{int}(e)$, and $V_x \cup V_y$ contains an open neighborhood of γ_{xy} .*

Proof. For $z \in \mathcal{S} - \{x, y\}$, let $\alpha_x \in [0, \pi]$ be the angle between z and x as measured from m , and let α_y be the angle from z to y . Since m is in the interior of the geodesic arc γ_{xy} we have $\alpha_x + \alpha_y = \pi$, so one of α_x and α_y is at most $\pi/2$. Assuming (without loss of generality) that $\alpha_x \leq \pi/2$, the hyperbolic law of cosines gives:

$$\begin{aligned} \cosh d(x, z) &= \cosh d(m, z) \cosh d(x, m) - \sinh d(m, z) \sinh d(x, m) \cos \alpha_x \\ &\leq \cosh d(m, z) \cosh d(x, m) \end{aligned}$$

Let $R_0 = B_0/2$. The “half-angle identity” for hyperbolic cosine implies that R_0 satisfies

$$\cosh R_0 = \sqrt{\frac{1}{2}(\cosh B_0 + 1)} = \sqrt{\cosh(2R)}$$


If $d(x, y) \leq B_0$ then $d(x, m) = \frac{1}{2}d(x, y) < R_0$. Since $d(x, z) \geq 2R$, combining expressions above yields:

$$\cosh d(m, z) \geq \frac{\cosh d(x, z)}{\cosh d(x, m)} \geq \frac{\cosh(2R)}{\sqrt{\cosh(2R)}} = \cosh R_0$$

Thus m is at least as close to x and y as to z and, since $z \in \mathcal{S}$ is arbitrary, $m \in V_x \cap V_y$.

If $d(x, y) < B_0$, let $\eta = \frac{\cosh R_0}{\cosh d(x, m)} > 1$. The inequality above gives $\cosh d(m, z) \geq \eta \cdot \cosh R_0$ in this case. Thus if $R_1 = \cosh^{-1}(\eta \cosh R_0)$ and $\delta = R_1 - R_0$, for $p \in B_{\delta/2}(m)$ the triangle inequality gives:

$$d(p, x) < R_0 + \delta/2 = R_1 - \delta/2 < d(p, z)$$


 FIGURE 2.1. Two possibilities from Lemma 2.6. Q is shaded.

Thus in this case $B_{\delta/2}(m) \subset V_x \cup V_y$, and if γ_{xy}^\perp is the perpendicular bisector to γ_{xy} then $\gamma_{xy}^\perp \cap B_{\delta/2}(m) \subset e = V_x \cap V_y$. In particular, $m = \gamma_{xy}^\perp \cap \gamma_{xy} \in \text{int}(e)$, and $\text{int}(V_x) \cup B_{\delta/2}(m) \cup \text{int}(V_y)$ is an open neighborhood of γ_{xy} in $V_x \cup V_y$. \square

Lemma 2.6. *Let P be a centered polygon in \mathbb{H}^2 with center $v \in \text{int } P$. For a vertex x of P , let $Q_x \subset P$ be the quadrilateral with vertices v , x , and the midpoints of the edges of P containing x . Then $P = \bigcup Q_x$, taken over all vertices of P . For $y \in P$, $y \in Q_x$ if and only if $d(y, x) \leq d(y, x')$ for each vertex x' of P .*

Proof. Let J be the radius of P — ie, the distance from v to the vertices of P — and let x_\pm be the vertices adjacent to x on ∂P . The geodesic arc e_+ that joins v to the midpoint m_+ of the edge γ_+ of P containing x and x_+ meets γ_+ at a right angle, since it is the fixed axis of a reflective involution of the isosceles triangle Δ_+ with vertices v , x , and x_+ . Thus e_+ is contained in the perpendicular bisector γ_+^\perp . The same holds true for the other edge e_- of Q_x containing v , and it follows that points of Q_x are at least as close to x as to either of x_\pm .

The center v is in Q_x and satisfies $d(v, x) = J = d(v, x')$ for all other vertices x' of P , so the conclusion holds for v . Fix $y \in Q_x - \{v\}$. If x' is a vertex of P other than x , x_+ or x_- , then the geodesic arc from y to x' crosses one of e_\pm , say e_+ . Let T be the triangle with vertices v , y , and x , and let T' have vertices v , y , and x' . Each of T and T' has an edge with length $d(y, v)$ and an edge with length $J = d(x, v) = d(x', v)$. We consider two cases.

If the geodesic arc from y to x' crosses the arc from v to x , as on the left-hand side of Figure 2.1, then clearly the angle θ' of T' at v is larger than the angle θ of T at v . Hence the hyperbolic law of cosines implies in this case that $d(x, y) < d(x', y)$.

If not, then $\theta \leq \theta_+/2$, where θ_+ is the angle at v of the isosceles triangle Δ_+ determined by v , x , and x_+ . This is because y lies in Δ_+ between its bisector e_+ and the edge joining v to x (see the right-hand side of Figure 2.1). On the other hand, the arc from y to x' exits Δ_+ at a point in the edge joining v to x_+ , since it crosses e_+ . Therefore $\theta' > \theta/2 \geq \theta$, and again by the hyperbolic law of cosines we have $d(x, y) < d(x', y)$.

For $y \neq v$, the geodesic ray from v in the direction of y exits P at a point in some edge γ . By construction $\gamma \subset Q_x \cup Q_{x'}$, where x and x' are the endpoints of γ . Thus y is in Q_x or $Q_{x'}$, say Q_x , since $v \in Q_x \cap Q_{x'}$ and these are convex quadrilaterals. For another vertex x'' ,

if $d(y, x'') \leq d(y, x)$ then $d(y, x'') = d(y, x)$, by the property of Q_x that we showed above. It follows that x is adjacent to x'' on ∂P , since we showed that $d(x'', y) > d(x, y)$ otherwise. In this case $Q_x \cap Q_{x''}$ is the intersection of Q_x with the equidistant locus of x and x'' by construction, so y is in $Q_x \cap Q_{x''}$. The lemma follows. \square

Lemma 2.7. *Let P be a centered polygon in \mathbb{H}^2 with radius J and center v . For each $y \in P - \{v\}$ there is a vertex x such that $d(x, y) < J$.*

Proof. Let $\{x_i\}_{i=0}^n$ be the set of vertices of P , enumerated so that for each i there is an edge of P containing x_i and x_{i+1} (with $i+1$ taken modulo n), and for each i let Q_i be the quadrilateral associated to x_i by the construction of Lemma 2.6. It is clear by construction that $P = \bigcup_{i=0}^{n-1} Q_i$, so for any $y \in P - \{v\}$ there exists $i \in \{0, \dots, n-1\}$ so that $y \in Q_i$. The conclusion of Lemma 2.6 implies that x_i is a closest vertex to y .

Claim 2.7.1. For any $i \in \{0, 1, \dots, n-1\}$ and $y \in Q_i - \{v\}$, $d(y, x_i) < J$.

Proof. The geodesic ray from x_i through y intersects one of the edges of Q_i containing v at a point y_0 . If $y_0 = v$, then since $y \neq v$ is on the geodesic arc joining x_i to v , the claim follows immediately. Otherwise let us consider the right triangle determined by x_i , y_0 , and the other endpoint of the edge containing v and y_0 , call it m . The hyperbolic law of cosines gives:

$$\cosh d(x_i, y_0) = \cosh d(x_i, m) \cosh d(m, y_0) = \cosh R \cosh d(m, y_0)$$

Since y_0 is contained in the geodesic arc joining m to v , and is not v , we have $d(m, y_0) < d(m, v)$. When applied to the triangle determined by x_i , v , and m , the hyperbolic law of cosines gives $\cosh R \cosh d(m, v) = \cosh J$. Thus by the above $d(x_i, y_0) < J$, and since $d(x_i, y) \leq d(x_i, y_0)$, the claim follows. \square

The lemma follows immediately. \square

Proposition 2.8. *Let $\mathcal{S} \subset \mathbb{H}^2$ have injectivity radius $R > 0$. If a centered n -gon P in \mathbb{H}^2 has vertices in \mathcal{S} , sides of length less than B_0 (from Lemma 2.5), and $\text{int}(P) \cap \mathcal{S} = \emptyset$, then:*

- (1) $P = P_v$ as in Lemma 1.9, where v is the center of P ; and
- (2) the Voronoi tessellation V determined by \mathcal{S} satisfies $V^{(0)} \cap P = \{v\}$, and $V^{(1)} \cap P$ is the union of geodesic arcs joining v to the midpoint of each side of P ; and

Proof. Let the vertices of P be cyclically ordered x_0, \dots, x_{n-1} in the sense of [3, Definition 1.3], and for each i let γ_i be the edge of P joining x_{i-1} to x_i (taking $i-1$ modulo n). By Lemma 2.5, γ_i has an open neighborhood contained in $V_{x_{i-1}} \cup V_{x_i}$ for each i . Since $\text{int } P \cap \mathcal{S} = \emptyset$ and Voronoi cells are connected, it follows that $P \subset \bigcup_{i=0}^{n-1} V_{x_i}$.

For each $i \in \{0, \dots, n-1\}$, Lemma 2.6 implies that the quadrilateral Q_{x_i} described there is contained in V_{x_i} , since its points are as close to x_i as any x_j for $j \neq i$. Since $P = \bigcup_{i=0}^{n-1} Q_i$ it follows that $Q_i = P \cap V_{x_i}$ for each i . The description of the Voronoi tessellation follows immediately; in particular, $Q_{x_{i-1}} \cap Q_{x_i}$ is contained in an edge $e_i = V_{x_{i-1}} \cap V_{x_i}$ containing v for each i . For each i , γ_i is the geometric dual to e_i , so Lemma 1.9 implies that $P = P_v$. \square

Corollary 2.9. *Let $\mathcal{S} \subset \mathbb{H}^2$ have injectivity radius $R > 0$. For $3 \leq n \leq 6$, a cyclic n -gon P with vertices in \mathcal{S} and all sides of length $2R$ satisfies the conclusions of Proposition 2.8.*

Proof. Since $B_0 > 2R$ by its definition in Lemma 2.5, the result will follow from Proposition 2.8 once we show that P is centered and $\text{int } P \cap \mathcal{S} = \emptyset$. P is represented by $(d, \dots, d) \in \widetilde{\mathcal{AC}}_n$. [3, Lemma 6.6] implies that $(d, \dots, d) \in \widetilde{\mathcal{C}}_n$, and hence that P is centered, and furthermore that its radius $J = J_n(d)$ satisfies $\sinh J = \sinh(d/2)/\sin(\pi/n)$. If $n \leq 6$ then $\sin(\pi/n) \geq 1/2$, so $\sinh J \leq 2 \sinh R < \sinh(2R)$. Since the hyperbolic sine is increasing on \mathbb{R}^+ it follows that $J < 2R$. Thus since the points of P have distance at most J from x by Lemma 2.7, and $J < 2R$ in this case, P has no points of \mathcal{S} but the x_i . \square

Corollary 2.10. *Let F_α and $x_\alpha \in F_\alpha$ be as in Example 2.2. The Delaunay tessellation of F_α determined by $\mathcal{S} = \{x_\alpha\}$ is the triangulation by the projections of T_1, \dots, T_6 described there, and each edge intersects the interior of its geometric dual.*

Proof. Let $O_\alpha = T_1 \cup \dots \cup T_6 \subset \mathbb{H}^2$, as described in Example 2.2. Given a scheme for isometrically pairing edges of O_α to produce F_α as in Example 2.2, for each pair of edges e and e' , there is an orientation-preserving isometry f of \mathbb{H}^2 with $f(e) = e'$ and $f(O_\alpha) \cap O_\alpha = e'$. The Poincaré polyhedron theorem asserts that the set of these edge pairings generates a discrete group Π of isometries with fundamental domain O_α , and that the quotient map $\mathbb{H}^2 \rightarrow \mathbb{H}^2/\Pi = F_\alpha$ is a locally isometric universal covering. In particular, Π -translates of O_α tessellate \mathbb{H}^2 .

Since O_α is itself tessellated by the T_i , \mathbb{H}^2 is tessellated by Π -translates of these six triangles. The preimage $\tilde{\mathcal{S}}$ of \mathcal{S} in \mathbb{H}^2 is the set of vertices of Π -translates of O_α , so the vertices of any Π -translate of any T_i are in $\tilde{\mathcal{S}}$. F_α has injectivity radius r_α at x_α , so $\tilde{\mathcal{S}}$ also has injectivity radius r_α . Thus since T_i has edge length $d_\alpha = 2r_\alpha$ for each i , Corollary 2.9 implies that each translate of each T_i is a two-cell of the Delaunay tessellation of \mathbb{H}^2 determined by $\tilde{\mathcal{S}}$. Lemma 2.5 further implies that each edge of each T_i intersects the interior of its geometric dual, and the conclusion for F_α follows from Definition 1.13. \square

The Voronoi tessellation from Example 2.2 is easily described. A similar proof establishes:

Corollary 2.11. *Let T_β , Q , F_β , and $x_\beta \in F_\beta$ be as in Example 2.4. The Delaunay tessellation of F_β determined by $\{x_\beta\}$ is the decomposition described there, into T_1, \dots, T_4 and Q , all with side lengths d_β . Each edge intersects the interior of its geometric dual.*

In particular, the Delaunay tessellation of F_β is not a triangulation. The example below shows that the conclusion of Theorem 0.1 fails upon slightly relaxing the hypothesis $r \geq r_\beta$.

Example 2.12. We will produce a family of surfaces F_t by perturbing the surface F_β from Example 2.4. By [3, Lemma 6.8], the equation determining d_β can be rewritten as

$$4\pi = 4 \cdot D_{0,3}(d_\beta) + 2 \cdot D_0(b_\beta, d_\beta, d_\beta),$$

where $b_\beta = \cosh^{-1}(2 \cosh d_\beta - 1) = b_0(d_\beta, d_\beta)$. This reflects the geometric observation, also in [3, Lemma 6.8], that the quadrilateral Q from Example 2.4 has a diagonal that contains

its center and divides it into triangles T_0 and T'_0 , each with side length collection $(b_\beta, d_\beta, d_\beta)$. Re-naming if necessary, we will assume that T'_0 shares an edge with T_1 from Example 2.4.

For t near 0 let $d_t = d_\beta + t$, $b_t = \cosh^{-1}(2 \cosh d_t - 1)$, and let $d_1(t)$ satisfy $d_1(0) = d_\beta$ and $f(t, d_1(t)) \equiv 4\pi$, where:

$$(2.12.1) \quad f(t, d) = 3 \cdot D_{0,3}(d_t) + D_0(d, d_t, d_t) + D_0(b_t, d_t, d) + D_0(b_t, d_t, d_t)$$

We note that $d_0 = d_\beta$ and $b_0 = b_\beta$, and by comparing with the equation above one finds that $f(0, d_\beta) = 4\pi$. To produce $d_1(t)$ we note that [3, Proposition 5.5] implies:

$$(2.12.2) \quad \frac{\partial f}{\partial d}(0, d_\beta) = \sqrt{\frac{1}{\cosh^2(d_\beta/2)} - \frac{1}{\cosh^2 J_3(d_\beta)}} > 0,$$

since $(d_\beta, d_\beta, d_\beta) \in \tilde{\mathcal{C}}_3$ by [3, Lemma 6.6], and $J_3(d_\beta) < d_\beta/2$. Therefore the implicit function theorem yields $\epsilon > 0$ and a function d_1 on $(-\epsilon, \epsilon)$ with $d_1(0) = d_\beta$ and $f(t, d_1(t)) = 4\pi$ for each $t \in (-\epsilon, \epsilon)$. With a computation analogous to the above it is possible to show that $\partial f / \partial t$ is also positive at $(0, d_\beta)$, and this further implies that d_1 decreases in t .

By [3, Lemma 6.6], $(d_t, d_t, d_t) \in \tilde{\mathcal{C}}_3$ for each t such that $d_t > 0$; we may as well assume this is all of $(-\epsilon, \epsilon)$. Moreover, $(b_t, d_t, d_t) \in \widetilde{\mathcal{BC}}_3$ for each such t by the definition of b_t and [3, Lemma 6.2]. We may also assume that for each $t \in (-\epsilon, \epsilon)$, each of $(b_t, d_t, d_1(t))$, and $(d_1(t), d_t, d_t)$ is in $\widetilde{\mathcal{AC}}_3$, since this set is open in \mathbb{R}^3 .

For each t , let $T_2(t)$, $T_3(t)$, and $T_4(t)$ be centered triangles with all side lengths d_t . Let $T_0(t)$ be a cyclic triangle with cyclically ordered side length collection (b_t, d_t, d_t) , let $T'_0(t)$ have side length collection $(b_t, d_t, d_1(t))$, and let $T_1(t)$ have side length collection $(d_1(t), d_t, d_t)$. That these exist follows from [3, Definition 3.1] and [3, Proposition 2.7].

Note that $T_i(0) = T_i$ for $0 \leq i \leq 4$, and $T'_0(0) = T'_0$. Arranging the triangles in \mathbb{H}^2 so that at time 0 their union is O_β , and they have the same combinatorial pattern of intersection for all time, their union at each time t is an octagon O_t with all side lengths d_t and area 4π by construction. The total angle defect of O_t is thus 2π , so an isometric edge-pairing scheme that is combinatorially identical to that for O_β produces a surface F_t . Let $x_t \in F_t$ be the quotient of the vertices of O_t . One can show as in the previous examples that for each $t < 0$, F_t has injectivity radius $r_t = d_t/2$ at x_t . Since d_1 decreases in t , $d_1(t) < d_t$ for $t > 0$, so F_t has injectivity radius at most $d_1(t)/2$ for such t .

We prove in the lemma below that for $t < 0$, the conclusion of Theorem 0.1 does not apply to the surfaces F_t from Example 2.12, though the Delaunay tessellation is a triangulation. Recall from the example that for such t , F_t has injectivity radius $r(t) < r_\beta$ at x_t .

Lemma 2.13. *For O_t , F_t , and $x_t \in F_t$ as in Example 2.12, and $t < 0$ but near to it, the triangulation that F_t inherits from O_t is its Delaunay tessellation determined by $\{x_t\}$. Each edge of this triangulation intersects the interior of its geometric dual edge except for $T_0(t) \cap T'_0(t)$, which intersects an endpoint of its geometric dual.*




FIGURE 2.2. $P = T_0 \cup T_1(\delta) \cup T_2(\delta)$ and its intersection with $\tilde{V}^{(1)}$.

Proof. For each t , the Poincaré polyhedron theorem implies that Π_t -translates of O_t tessellate \mathbb{H}^2 , where Π_t is the group generated by the edge-pairing isometries of O_t yielding F_t . It follows that \mathbb{H}^2 is triangulated by Π_t -translates of $T'_0(t)$ and the $T_i(t)$ for $0 \leq i \leq 4$. The preimage $\tilde{\mathcal{S}}_t$ of x_t in \mathbb{H}^2 is the set of vertices of Π_t -translates of O_t .

For $i = 2, 3$, or 4 , since $T_i(t)$ has all side lengths equal to $d_t = 2r_t$, Corollary 2.9 implies that each of its Π_t -translates is a 2-cell of the Delaunay tessellation \tilde{P} of \mathbb{H}^2 determined by $\tilde{\mathcal{S}}$. Lemma 2.5 applies to the edges of these polygons, and asserting in particular that each intersects the interior of its geometric dual.

For reference we have depicted $T_0(t) \cup T'_0(t) \cup T_1(t)$ in Figure 2.2, and labeled its vertices. By construction, each x_i is in $\tilde{\mathcal{S}}$, and each frontier edge has length d_t . Lemma 2.5 thus implies that each frontier edge is in $\tilde{P}^{(1)}$, intersects the interior of its geometric dual, and has an open neighborhood contained in the union of the Voronoi cells determined by its endpoints.

The edge $T'_0(t) \cap T_1(t)$, joining x_0 to x_3 in the figure, has length $d_1(t)$. Since $d_1(0) = d_\beta < b_\beta = b_0$, the inequality $d_1(t) < b_t$ holds for t near to 0. For such t , since $d_1(t) < b_t = \cosh^{-1}(2 \cosh d_t - 1)$, Lemma 2.5 asserts that $T'_0(t) \cap T_1(t)$ is in $\tilde{P}^{(1)}$ and intersects the interior of its geometric dual, and [3, Lemma 6.2] implies that $(d_1(t), d_t, d_t) \in \tilde{\mathcal{C}}_3$. Therefore $T_1(t)$ is centered, and since it intersects $\tilde{\mathcal{S}}$ in its vertex set, Proposition 2.8 implies it is a Delaunay 2-cell intersecting $\tilde{V}^{(1)}$ as illustrated in Figure 2.2.

It remains to consider $T_0(t) \cup T'_0(t)$. We have already showed that its frontier is in $P^{(1)}$, and moreover has an open neighborhood contained in $V_{x_0} \cup V_{x_1} \cup V_{x_2} \cup V_{x_3}$. Since $T_0(t) \cup T'_0(t) \cap \tilde{\mathcal{S}} = \{x_0, x_1, x_2, x_3\}$ and Voronoi cells are connected, $T_0(t) \cup T'_0(t)$ is entirely contained in $V_{x_0} \cup V_{x_1} \cup V_{x_2} \cup V_{x_3}$. We claim that V_{x_0} intersects V_{x_2} in an edge e , whose geometric dual $T_0(t) \cap T'_0(t)$ is thus in $\tilde{P}^{(1)}$ and furthermore intersects e in an endpoint. It will follow immediately that $T_0(t)$ and $T'_0(t)$ are each Delaunay 2-cells.

Since $\cosh b_t = 2 \cosh d_t - 1$ by construction, [3, Lemma 6.2] implies that $(b_t, d_t, d_t) \in \tilde{\mathcal{BC}}_3$. Therefore $T_0(t)$ has its center v_0 at the midpoint of its longest edge $T_0(t) \cap T'_0(t)$ by [3, Lemma 3.9]. On the other hand, we pointed out below (2.12.2) that d_1 is decreasing in t , so $d_1(t) > d_1(0) = d_\beta > d_t$ for $t < 0$. Therefore $\cosh b_t < \cosh d_t + \cosh d_1(t) - 1$, so $(b_t, d_t, d_1(t)) \in \tilde{\mathcal{C}}_3$ by [3, Lemma 6.2] again. Thus $T'_0(t)$ has its center v_1 in its interior.

By definition v_0 is equidistant from x_0, x_1 and x_2 , and v_1 is equidistant from x_0, x_2 , and x_3 . The quadrilaterals Q_{x_0} and Q_{x_2} in $T'_0(t)$ supplied by Lemma 2.6 contain v_0 in their intersection, so since $v_0 \neq v_1$ it is not contained in Q_{x_3} . Lemma 2.6 thus implies that $d(v_0, x_3) > d(v_0, x_0)$, and hence that $v_0 = V_{x_0} \cap V_{x_1} \cap V_{x_2}$.

Since $T_0(t)$ is isosceles, the geodesic arc γ from x_1 to v_0 intersects $T_0(t) \cap T'_0(t)$ at a right angle. Since the geodesic arc γ' from v_1 to v_0 also meets $T_0(t) \cap T'_0(t)$, their union is geodesic. It follows that $d(x_1, v_1) = d(x_1, v_0) + d(v_0, v_1)$. On the other hand, $d(x_2, v_1)$ satisfies

$$\cosh d(x_2, v_1) = \cosh d(x_2, v_0) \cosh d(v_0, v_1) = \cosh d(x_1, v_0) \cosh d(v_0, v_1)$$

by the hyperbolic law of cosines. The angle addition formula for hyperbolic cosine therefore implies that $d(x_1, v_1) > d(x_2, v_1)$, so $v_1 = V_{x_0} \cap V_{x_2} \cap V_{x_3}$, and it follows that $V_{x_0} \cap V_{x_2}$ contains an edge joining v_0 to v_1 . The lemma follows. \square

Remark 2.14. The construction of Example 2.12 may be modified, by increasing d_0 and reducing δ , to produce deformations of F_β in which an edge of the Delaunay tessellation does not intersect its geometric dual at all.

3. THE CENTERED DUAL TO THE VORONOI TESSELLATION

Our task in this section is to understand the “pathology” described in Lemma 2.13, in which an edge of V does not intersect the interior of its geometric dual. We will say that such an edge of V is “non-centered,” and relate (non-)centeredness of edges to (non-)centeredness of vertex polygons in Lemma 3.3. The set of non-centered edges has restricted combinatorics: its components are sub-trees of $P^{(1)}$, each with a canonical root vertex (Lemma 3.6). We organize the Delaunay polygons corresponding to vertices of such a component into a 2-cell of the “centered dual decomposition” P_c , in Definition 3.18.

Definition 3.1. Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete. We will say an edge e of V is *centered* if e intersects its geometric dual γ_{xy} at a point in $\text{int } e$. If e is not centered, orient it pointing away from γ_{xy} .

If V is the Voronoi tessellation of a closed surface F determined by a finite set \mathcal{S} , we say an edge e of V is *centered* if and only if one (and hence all) of its lifts to \tilde{V} is centered, where $\tilde{V} = \pi^{-1}(V) \subset \mathbb{H}^2$ is the Voronoi tessellation of \mathbb{H}^2 determined by $\tilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$. If e is not centered, let it inherit an orientation from a lift \tilde{e} .

As indicated above, the action of $\pi_1 F$ on $\tilde{V}^{(1)}$ preserves (non-)centeredness of edges, and also the orientation of non-centered edges.

Fact. Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete. For $x \in \mathcal{S}$, an edge e of V_x is non-centered with initial vertex v if and only if the angle α at v , measured in V_x between e and the geodesic segment joining v to x , is at least $\pi/2$.

This is because there is a right triangle with vertices at x and v and edges contained in γ_{xy} and γ_{xy}^\perp , where γ_{xy} is the geometric dual to e ; ie, $e = V_x \cap V_y$. This triangle has angle equal to either α or $\pi - \alpha$ at v , depending on the case above; see Figure 3.1.




FIGURE 3.1. Centered and non-centered edges.

If w is the other endpoint of e then since $x \in P_v \cap P_w$, the fact above and the hyperbolic law of cosines imply:

$$(3.1.1) \quad \cosh J_w = \cosh \ell(e) \cosh J_v - \sinh \ell(e) \sinh J_v \cos \alpha$$

Because $\cos \alpha \leq 0$ if $\alpha \geq \pi/2$, we have:

Lemma 3.2. *Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete. For $x \in \mathcal{S}$, if e is a non-centered edge of V_x oriented as prescribed in Definition 3.1, with initial vertex v and terminal vertex w , then $J_v < J_w$.*

Below we relate centeredness of edges of V to that of 2-cells of the Delaunay tessellation.

Lemma 3.3. *Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete. For $v \in V^{(0)}$, P_v is non-centered if and only if v is the initial vertex of a non-centered edge e of V . If this is so, the geometric dual γ to e is the unique longest edge of P_v , and $P \cup T(e, v)$ is a convex polygon, where $T(e, v)$ is the triangle determined by v and $\partial\gamma$.*

Proof. Suppose first that v is the initial vertex of a non-centered edge $e = V_x \cap V_y$, and let \mathcal{H}' be the half-space containing e and bounded by the geodesic containing x and y . The circle C with radius J_v and center v intersects $\partial\mathcal{H}'$ in $\{x, y\}$. If α is the angle at v between e and the geodesic arc to x , then by the Fact above, $\alpha \geq \pi/2$. The hyperbolic law of cosines implies that $z \in C$ is in the interior of \mathcal{H}' if and only if the angle α' at v between e and the geodesic arc to z is less than α . Thus if w is the other endpoint of e for such z :

$$\cosh d(z, w) = \cosh \ell(e) \cosh J_v - \sinh \ell(e) \sinh J_v \cos \alpha'$$

Since $\alpha' < \alpha$, comparing with (3.1.1) we find that $d(z, w) < J_w$, so the intersection of C with the interior of \mathcal{H} is entirely contained in $B_{J_w}(w)$. Therefore by Lemma 1.2 it contains no points of \mathcal{S} . Since all vertices of P_v are on C , it follows that P_v is contained in the half-plane \mathcal{H} opposite \mathcal{H}' , and hence that $v \notin \text{int } P_v$. Thus P_v is non-centered by [3, Definition 1.1].

Assume now that P_v is not centered and apply [3, Lemma 1.5]. This produces an edge γ of P_v and a half-space \mathcal{H} containing P_v and bounded by the geodesic containing γ , such that v is in the half-space \mathcal{H}' opposite \mathcal{H} . [3, Lemma 1.5] further asserts that $P \cup T(e, v)$ is a convex polygon; also, γ is the unique longest edge of P_v , by [3, Corollary 1.11]. We claim that the other endpoint w of the geometric dual e to γ is further from \mathcal{H} than v , and hence that e is non-centered with initial vertex v .

If w is closer to \mathcal{H} than v (this includes the possibility $w \in \mathcal{H}$), then e intersects the geodesic joining v to x in an angle of $\alpha \leq \pi/2$. J_w again satisfies (3.1.1), and if C is the circle of radius J_v centered at v , the hyperbolic law of cosines again implies that $z \in C$ is in $\text{int } \mathcal{H}$ if and only if the angle at v between e and the geodesic joining v to z is less than α . As in the previous case, this implies that the distance from the other vertices of P_v to w is less than J_w , contradicting Lemma 1.2. Therefore w is further from \mathcal{H} than v . \square

If $v \in V^{(0)}$ is the initial vertex of a non-centered edge e , the fact that the geometric dual to e is the *unique* longest edge of P_v immediately implies the following.

Corollary 3.4. *Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete. No $v \in V^{(0)}$ is the initial vertex of more than one non-centered edge.*

Definition 3.5. If V is the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete, let $V_n^{(1)} \subset V^{(1)}$ be the union of the non-centered edges. If F is a closed surface, we define $V_n^{(1)}$ in the same way for the Voronoi tessellation V determined by $\mathcal{S} \subset F$ finite.

Below, given a graph G we will say that $\gamma = e_0 \cup e_1 \cup \dots \cup e_{n-1}$ is an *edge path* if e_i is an edge of G for each i and $e_i \cap e_{i-1} \neq \emptyset$ for $i > 0$. An edge path γ as above is *reduced* if $e_i \neq e_{i-1}$ for each $i > 0$, and γ is *closed* if $e_0 \cap e_{n-1} \neq \emptyset$.

Lemma 3.6. *Let \tilde{V} be the Voronoi tessellation determined by $\tilde{\mathcal{S}} \subset \mathbb{H}^2$ closed and discrete. Each component T of $\tilde{V}_n^{(1)}$ is a tree. If $\tilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$, where $\pi: \mathbb{H}^2 \rightarrow F$ is the universal cover to a closed surface F , and $\mathcal{S} \subset F$ is finite, then T is finite, with a unique vertex v_T such that $J_{v_T} > J_v$ for all $v \in T^{(0)} - \{v_T\}$, and T projects homeomorphically to F .*

Proof. Suppose that a component T of $\tilde{V}_n^{(1)}$ admits closed, reduced edge paths, and let $\gamma = e_0 \cup e_1 \cup \dots \cup e_{n-1}$ be a shortest such. Orienting the e_i as in Definition 3.1, we may assume (after re-numbering if necessary) that e_0 points toward $e_0 \cap e_{n-1}$. We claim that then e_i points to $e_i \cap e_{i-1}$ for each $i > 0$ as well. Otherwise, for the minimal $i > 0$ such that e_i points toward e_{i+1} it would follow that the vertex $e_i \cap e_{i-1}$ was the initial vertex of both e_i and e_{i-1} , contradicting Corollary 3.4.

Let $v_0 = e_0 \cap e_{n-1} \in V^{(0)}$, and for $i > 1$ take $v_i = e_i \cap e_{i-1}$. Applying Lemma 3.2 to e_i for each i , we find that $J_{v_i} > J_{v_{i+1}}$. By induction this gives $J_{v_0} > J_{v_{n-1}}$; but since e_{n-1} points to v_{n-1} Lemma 3.2 implies that $J_{v_{n-1}}$ must exceed J_{v_0} , a contradiction. Thus no component of $\tilde{V}_n^{(1)}$ admits closed, reduced edge paths, so each is a tree.

If $\tilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$ as in the hypotheses then the set $\{J_v \mid v \in T^{(0)}\}$ has only finitely many distinct elements, since $J_v = J_{v'}$ if v and v' project to the same point of \mathcal{S} . Thus take $v_T \in T^{(0)}$ with J_{v_T} maximal. We claim that $J_v < J_{v_T}$ for each $v \in T^{(0)} - \{v_T\}$.

If there exists $v \in T^{(0)} - \{v_T\}$ with $J_v = J_{v_T}$, let $\gamma = e_0 \cup \dots \cup e_{n-1}$ be a reduced edge path joining v_T to v . We may assume that v_T is the endpoint of e_0 not in e_1 , and $v \in e_{n-1} - e_{n-2}$. Lemma 3.2 implies that e_0 points toward v_T and e_{n-1} towards v . Thus if $i > 0$ is minimal such that e_i does not point toward e_{i-1} , $v_i = e_i \cap e_{i-1}$ is the initial endpoint of e_i and e_{i-1} , contradicting Corollary 3.4. This proves the claim.

Since covering transformations exchange components of $\tilde{V}_n^{(1)}$, if $\gamma.T \cap T \neq \emptyset$ for some $\gamma \in \pi_1 F - \{1\}$ then $\gamma.T = T$. Since $J_{\gamma.v} = J_v$ for each $v \in T^{(0)}$, the claim above would imply that $\gamma.v_T = v_T$ for such γ , contradicting freeness of the $\pi_1 F$ -action. Therefore T does not intersect its $\pi_1 F$ -translates and thus projects homeomorphically to F . It follows that $T^{(0)}$, and hence also T , is finite. \square

It is a basic fact that any two distinct points in a tree are joined by a unique reduced edge path, and that each such path is homeomorphic to an embedded interval.

Lemma 3.7. *Let V be the Voronoi tessellation of a closed hyperbolic surface F determined by $\mathcal{S} \subset F$ finite, and let T be a component of $V_n^{(1)}$. For $v_T \in T^{(0)}$ as in Lemma 3.6, P_{v_T} is centered. For $v \in T^{(0)} - \{v_T\}$, P_v is not centered, and the reduced edge path joining v to v_T inherits an orientation from each of its constituent edges, pointing from v to v_T .*

Proof. Since v_T has maximal radius among $v \in T^{(0)}$, Lemma 3.2 implies that it is the terminal endpoint of each edge of T containing it. Since every other edge of V containing v_T is centered, Lemma 3.3 implies that P_{v_T} is centered.

For $v \in T^{(0)} - \{v_T\}$, let $e_0 \cup \dots \cup e_{n-1}$ be the reduced edge path joining v to v_T , take $v_i = e_i \cap e_{i-1}$ for $1 \leq i \leq n-1$, and let v_0 and v_n be the endpoints of e_0 and e_{n-1} not equal to v_1 and v_{n-1} , respectively. Re-numbering if necessary, we may assume that $v_0 = v$ and $v_n = v_T$. Then e_{n-1} has terminal endpoint v_n . If e_i does not point toward v_{i+1} for some $i < n-1$ then for the maximal such i , v_i is the initial vertex of e_i and e_{i+1} , contradicting Corollary 3.4. The e_i thus agree with the orientation on the edge path that points toward v_T . In particular, v is the initial vertex of e_0 , so P_v is not centered by Lemma 3.3. \square

Definition 3.8. Let V be the Voronoi tessellation of a closed hyperbolic surface F determined by $\mathcal{S} \subset F$ finite. Define the *centered dual graph* $P_c^{(1)}$ to $V^{(1)}$ as:

$$P_c^{(1)} = \bigcup \{ \gamma_{xy} \mid \gamma_{xy} \text{ is the geometric dual to a centered edge } e \subset V^{(1)} \} \subset P^{(1)}$$

Let $\tilde{P}_c^{(1)} \subset \tilde{P}^{(1)}$ be the preimage of $P_c^{(1)}$ in the universal cover.

It is easy to see that $\tilde{P}_c^{(1)}$ is the union of geometric duals to centered edges of \tilde{V} (cf. Definition 3.1). The centered dual graph has the structure of a subgraph of the one-skeleton $P^{(1)}$ of the Delaunay tessellation. It exhibits the behavior that one expects from a dual:

Lemma 3.9. *Let V be the Voronoi tessellation of a closed hyperbolic surface F determined by a finite set $\mathcal{S} \subset F$. If e is a centered edge of V and $\gamma \subset P_c^{(1)}$ its geometric dual, then $\gamma \cap V^{(1)} = \gamma \cap e$ is a single point. Furthermore, $e \cap P_c^{(1)} = e \cap \gamma$.*

Proof. Let \tilde{e} be a lift of e to \mathbb{H}^2 , and let x and $y \in \tilde{\mathcal{S}}$ be such that $\tilde{e} = V_x \cap V_y$. Then the geodesic arc γ_{xy} joining x to y projects to γ . Let $p = e \cap \gamma_{xy}$, and let $[x, p]$ and $[y, p]$ be the sub-arcs of γ_{xy} joining x and y , respectively, to p . Since V_x is convex with x in its interior, it contains $[x, p]$, and $[x, p] \cap \partial V_x = \{p\}$. The analogous assertion holds for V_y and $[y, p]$, and so $\gamma_{xy} \cap \tilde{V}^{(1)} = \{p\}$. Since this holds for any lift of e , the first claim follows.

The second claim follows from the first, since the set of centered edges of $V^{(1)}$ is in bijective correspondence with the edge set of $P_c^{(1)}$ by associating a centered edge to its dual. \square

Figure 3.1 shows that the conclusion of Lemma 3.9 does not hold for non-centered edges.

Lemma 3.6 implies that each cell V_x of V has at least one centered edge e , for otherwise some component of $V_n^{(1)}$ would contain the closed loop ∂V_x . Since the geometric dual of such an edge e is of the form γ_{xy} , it follows that the vertex set of $P_c^{(1)}$ is all of \mathcal{S} .

The interior of each Voronoi cell V_x is isometric to the interior of a compact, convex polygon in \mathbb{H}^2 . Therefore there is a “geometric” deformation retract $V_x - \{x\} \rightarrow \partial V_x \subset V^{(1)}$ along geodesic arcs connecting x to points on ∂V_x . Since $P_c^{(1)}$ intersects each V_x in a collection of such arcs, we have:

Lemma 3.10. *Let V be the Voronoi tessellation of a closed hyperbolic surface F determined by $\mathcal{S} \subset F$ finite, and let $P_c^{(1)}$ be the centered dual graph to $V^{(1)}$. There is a deformation retract $\rho_{\mathcal{S}}: F - P_c^{(1)} \rightarrow V^{(1)} - (V^{(1)} \cap P_c^{(1)})$ such that for each $x \in \mathcal{S}$, $\rho_{\mathcal{S}}$ restricts on $V_x - P_c^{(1)}$ to the restriction of the corresponding geometric deformation retract.*

Since $\rho_{\mathcal{S}}$ is a deformation retract, it determines a one-to-one correspondence between the set of components of $F - P_c^{(1)}$ and the set of components of $V^{(1)} - (V^{(1)} \cap \Gamma)$. We use this to give an initial description of the components of $F - \Gamma$. It will be convenient to first introduce another definition.

Definition 3.11. If V is a graph and T a subgraph, we define the *frontier* \mathcal{F} of T in V to be the set of pairs (e, v) where e is an edge of V that is not in T and $v \in e \cap T$.

We may refer just to “an edge” of the frontier of T , without reference to its vertices, but note that \mathcal{F} has two elements for each e not in T with both endpoints in T .

Lemma 3.12. *Let V be the Voronoi tessellation of a closed hyperbolic surface F determined by $\mathcal{S} \subset F$ finite, and let $P_c^{(1)}$ be the centered dual graph to $V^{(1)}$. Each component U of $F - P_c^{(1)}$ is homeomorphic to an open disk, and either:*

- (1) *there is a unique $v \in V^{(0)} \cap U$, each edge of V containing v is centered, and the universal cover maps $P_{\tilde{v}}$ to \overline{U} for (any) $\tilde{v} \in \pi^{-1}(v)$; or*
- (2) *U contains a unique component T of $V_n^{(1)}$, and $V^{(0)} \cap U = T^{(0)}$.*

Proof. Suppose for $v \in V^{(0)} \cap U$ that every edge of V containing v is centered. The same holds for $\tilde{v} \in \pi^{-1}(v)$, so Lemma 3.3 implies that the vertex polygon $P_{\tilde{v}}$ is centered, and hence that $\tilde{v} \in \text{int } P_{\tilde{v}}$. By Lemma 1.9, each edge of $P_{\tilde{v}}$ is the geometric dual to an edge containing \tilde{v} . Thus $\partial P_{\tilde{v}} \subset \tilde{P}_c^{(1)}$, so $\text{int } P_{\tilde{v}}$ is a component of $\mathbb{H}^2 - \tilde{P}_c^{(1)}$ containing \tilde{v} . Since π maps $\partial P_{\tilde{v}}$ into $P_c^{(1)}$ and $\text{int } P_{\tilde{v}}$ homeomorphically, $U = \pi(\text{int } P_{\tilde{v}})$.

In particular, since $\text{int } P_{\tilde{v}}$ is homeomorphic to an open disk, the same holds for U . Let the edges of \tilde{V} containing \tilde{v} be cyclically ordered e_0, \dots, e_{n-1} as in Lemma 1.9, with geometric dual γ_i for each i , and let $m_i = e_i \cap \gamma_i \in \text{int } e_i$. For each the quadrilateral Q_{x_i} constructed

in Lemma 2.6 is contained in V_{x_i} since its vertices are (cf. Definition 1.5), so the conclusion there implies that $P_{\tilde{v}} \subset \bigcup_{i=0}^{n-1} Q_{x_i} \subset \bigcup_{i=0}^{n-1} V_{x_i}$. Since $Q_{x_i} \cap \partial V_{x_i} = (e_i \cap P_{\tilde{v}}) \cup (e_{i+1} \cap P_{\tilde{v}})$ it follows that $v = V^{(0)} \cap U$.

If $v \in V^{(0)} \cap U$ is a vertex of a non-centered edge, then U contains the entire component T of $V_n^{(1)}$ containing v , since T does not intersect $P_c^{(1)}$. Let \tilde{T} be a lift of T to the universal cover. For each (e, v) in the frontier \mathcal{F} of \tilde{T} , e is centered, so its geometric dual γ intersects it and lies in $\tilde{P}_c^{(1)}$. For $v \in e \cap \tilde{T}$, let $[v, e \cap \gamma)$ refer to the component of $e - \gamma$ containing v . Then

$$S_{\tilde{T}} = T \cup \left(\bigcup \{[v, e \cap \gamma) \mid (e, v) \in \mathcal{F}\} \right)$$

is a connected open subset of $\tilde{V}^{(1)} - \tilde{P}_c^{(1)}$ with frontier in $\widetilde{P_c^{(1)}}$, and hence is a component of $\tilde{V}^{(1)} - \tilde{P}_c^{(1)}$. Again we find that by construction, $S_{\tilde{T}}$ deformation retracts to \tilde{T} (thus in particular is simply connected) and that $S_{\tilde{T}} \cap \tilde{V}^{(0)} = \tilde{T}^{(0)}$. Therefore $S_{\tilde{T}}$ projects homeomorphically to the component S_T of $V^{(1)} - P_c^{(1)}$ containing T .

Since S_T is a component of $V^{(1)} - P_c^{(1)}$ contained in U , and U is a component of $F - P_c^{(1)}$, $U = \rho_S^{-1}(S_T)$; in particular, $T^{(0)} = S_T \cap V^{(0)} = U \cap V^{(0)}$. Since ρ_S is a deformation retract, U is simply connected and hence lifts homeomorphically to the component \tilde{U} of $\mathbb{H}^2 - \tilde{P}_c^{(1)}$ containing $S_{\tilde{T}}$. This is homeomorphic to a disk by, say, the Riemann mapping theorem, and therefore so is U . \square

To better understand the structure of complementary components to $P_c^{(1)}$ that contain points of $V_n^{(1)}$, we introduce a new tool.

Definition 3.13. Let V be the Voronoi tessellation of \mathbb{H}^2 determined by $\mathcal{S} \subset F$ closed and discrete. For $v \in V^{(0)}$, an edge e of V containing v , and $x, y \in \mathcal{S}$ such that $e = V_x \cap V_y$, let $T(e, v)$ be the isosceles triangle with vertices v, x and y .

If V is the Voronoi tessellation of a closed hyperbolic surface F determined by $\mathcal{S} \subset F$ finite, let $T(e, v) = \pi(T(e, \tilde{v}))$, where $\tilde{v} \in \pi^{-1}(v)$.

The edges of $T(e, v)$ that join v to x and y , respectively, each have length J_v , and the third edge of (e, v) is the geometric dual γ_{xy} to e . If v and w are opposite endpoints of e , then $T(e, v)$ and $T(e, w)$ share the edge γ_{xy} . Whether their intersection is larger than this depends on whether e is centered — see Figure 3.2. In particular:

Lemma 3.14. *Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete. If e is a non-centered edge of V with initial vertex v and terminal vertex w , then $T(e, v) \subset T(e, w)$ and $T(e, v) \cap \partial T(e, w)$ is the geometric dual of e .*

Proof. Since e is non-centered it is contained on one side of the geodesic in \mathbb{H}^2 containing its geometric dual γ , so the nearer vertex v on e to γ is in the interior of $T(e, w)$. The result now follows from convexity. \square

Lemma 3.15. *Let V be the Voronoi tessellation determined by $\mathcal{S} \subset \mathbb{H}^2$ closed and discrete. For $v \in V^{(0)}$:*




FIGURE 3.2. Triangles $T(e, \tilde{v})$ and $T(e, \tilde{w})$ when e is centered (on the left) and not centered.

- (1) If P_v is centered then $P_v = \bigcup \{T(e, v) \mid v \in e\}$.
- (2) Otherwise, $P_v \cap T(e_v, v) = \gamma_v$ and $P_v \cup T(e_v, v) = \bigcup \{T(e', v) \mid e \neq e', v \in e'\}$, where e_v is the non-centered edge of V with initial vertex v and γ_v is its geometric dual.

Let γ and γ' be the respective geometric duals to e and $e' \neq e$ containing v . In case (1), $T(e, v) \cap T(e', v) = \{v\}$ if $\gamma \cap \gamma' = \emptyset$, and otherwise is an edge joining v to $\gamma \cap \gamma'$. This holds in case (2) for $e, e' \neq e_v$.

Proof. If the edges of P_v are enumerated $\gamma_0, \dots, \gamma_i$, for each i the triangle T_i described in the hypothesis of [3, Lemma 1.6] is identical to $T(e_i, v)$, where e_i is the geometric dual to γ_i . If P_v is non-centered then Lemma 3.3 implies that γ_v as defined above is its unique longest edge, so [3, Corollary 1.11] and [3, Lemma 1.5] imply that $T(e_v, v) \cap P_v = \gamma_v$. The decompositions of P_v and $P_v \cup T(e, v)$ described above follow directly from [3, Lemma 1.6]. \square

Proposition 3.16. *Let V be the Voronoi decomposition of a closed hyperbolic surface F determined by $\mathcal{S} \subset F$ finite, and let $P_c^{(1)}$ be the centered dual graph to $V^{(1)}$. For each component U of $F - P_c^{(1)}$, $\bar{U} = \bigcup_{v \in Q \cap V^{(0)}} P_v$.*

Proof. In case (1) of Lemma 3.12, the desired conclusion is proved there, so we will assume that U contains a component T of $V_n^{(1)}$. Let \tilde{T} be a lift of T to \mathbb{H}^2 , let \tilde{U} be the component of $\mathbb{H}^2 - \tilde{\Gamma}$ containing \tilde{T} , and let \tilde{Q} be the closure of \tilde{U} in \mathbb{H}^2 . Let \mathcal{F} be the frontier of \tilde{T} in $\tilde{V}^{(1)}$. For $(e, v) \in \mathcal{F}$, we claim that $T(e, v)$ is contained in \tilde{Q} .

Let x and $y \in \tilde{\mathcal{S}}$ be such that $e = V_x \cap V_y$. Then the side of $T(e, v)$ opposite v is $\gamma_{xy} \subset \tilde{P}_c^{(1)}$, since e is centered, and $T(e, v)$ is equal to the union of its intersections with V_x and V_y . Let $[v, \gamma_{xy}]$ be the sub-arc of e running from v to $e \cap \gamma_{xy}$. Lemma 3.9 implies that $[v, \gamma_{xy}] \cap \tilde{P}_c^{(1)} = \{e \cap \gamma_{xy}\}$, so since $v \in \tilde{T} \subset \tilde{U}$ it follows that $[v, \gamma_{xy}] \subset \tilde{Q}$.

$T(e, v) \cap V_x$ is the union of geodesic arcs in V_x joining x to points on $[v, \gamma_{xy}]$. Since $\tilde{P}_c^{(1)} \cap V_x$ is a union of geodesic arcs joining x to points of $\tilde{P}_c^{(1)} \cap \partial V_x$, and the only such point in $[v, \gamma_{xy}]$ is $e \cap \gamma_{xy}$, it follows that $T(e, v) \cap V_x$ intersects $\tilde{P}_c^{(1)}$ only in $\gamma_{xy} \cap V_x$. Since \tilde{U} contains a neighborhood of v and has its frontier in $\tilde{P}_c^{(1)}$, this implies that $T(e, v) \cap V_x \subset \tilde{Q}$. The analogous argument gives the analogous result for $T(e, v) \cap V_y$, and the claim is proved.

Lemma 3.15 implies that the interior of P_v intersects that of $T(e, v)$, so the claim above implies that the interior of P_v intersects \tilde{U} . By Lemma 1.9, the interior of P_v is a component of $\mathbb{H}^2 - \tilde{P}^{(1)}$. Therefore since $\tilde{P}_c^{(1)} \subset \tilde{P}^{(1)}$, $\text{int } P_v \subset \tilde{U}$, and hence $P_v \subset \tilde{Q}$.

If v and w are adjacent in \tilde{T} , then the corresponding vertex polygons P_v and P_w share an edge γ of \tilde{P} , the geometric dual to the edge of T joining v to w . Since this edge is non-centered, γ intersects $\tilde{P}_c^{(1)}$ only at its endpoints. Since $\text{int } P_v$ and $\text{int } P_w$ are each components of $\mathbb{H}^2 - \tilde{P}_c^{(1)}$, it follows that $P_v \subset \tilde{Q}$ if and only if $P_w \subset \tilde{Q}$.

We have already proved that $P_v \subset \tilde{Q}$ for any $v \in \tilde{T}^{(0)}$ such that $(e, v) \in \mathcal{F}$ for some edge e , so since \tilde{T} is connected, the previous paragraph and an inductive argument show that $\bigcup_{v \in \tilde{T}^{(0)}} P_v \subset \tilde{Q}$. Projecting to F it follows that $\bigcup_{v \in T^{(0)}} P_v \subset \bar{U}$.

It remains to show that $\bigcup_{v \in T^{(0)}} P_v$ is not properly contained in \bar{U} . If it were, then there would exist $v' \notin T^{(0)}$ such that $P_{v'} \subset \bar{U}$. But then Lemma 3.12 implies that v' is contained in a different component U' of $F - P_c^{(1)}$, so by the above $P_{v'} \subset \bar{U}'$. But since $P_{v'}$ has non-empty interior, this is a contradiction. \square

Corollary 3.17. *Let V be the Voronoi decomposition of a closed hyperbolic surface F determined by $\mathcal{S} \subset F$ finite, and let $P_c^{(1)}$ be the centered dual graph to $V^{(1)}$. For each component U of $F - P_c^{(1)}$, the completion \bar{Q} of the induced path metric on U is homeomorphic to a closed disk. If U contains a component T of $V_n^{(1)}$ with frontier \mathcal{F} in $V^{(1)}$, then:*

$$\bar{Q} - U = \bigcup \{ \gamma_{(e,v)} \mid (e, v) \in \mathcal{F} \},$$

where $\gamma_{(e,v)}$ is isometric to the geometric dual to e for each $(e, v) \in \mathcal{F}$.

A brief proof sketch: \bar{Q} is homeomorphic to the complement in U of a small neighborhood of its frontier, itself a closed disk. If e is an edge of V that is not in T but has both endpoints in it, then its geometric dual γ contributes two edges to \bar{Q} — one for each side — but only one to the closure \bar{Q} of U . This is why we use the induced path metric. It holds even in \mathbb{H}^2 : a lift of U to $\tilde{U} \subset \mathbb{H}^2 - \tilde{P}_c^{(1)}$ determines a map from \bar{Q} to the closure $\tilde{\bar{Q}}$ of \tilde{U} that is two-to-one over each edge of $\tilde{P}_c^{(1)}$ geometrically dual to a lift of e as above, and injective elsewhere.

Definition 3.18. Let V be the Voronoi tessellation of a closed hyperbolic surface F determined by $\mathcal{S} \subset F$ finite. Define the *centered dual decomposition* P_c of F to be the cell complex with $P_c^{(0)} = \mathcal{S}$, $P_c^{(1)}$ as described in Definition 3.8, and 2-cells as in Corollary 3.17.

If U is a component of $F - P_c^{(1)}$ containing a component T of $V_n^{(1)}$, we will refer to its closure \bar{Q} as a 2-cell of P_c with *vertex set* $\bar{Q} \cap \mathcal{S} = \bigcup_{v \in T^{(0)}} P_v \cap T^{(0)}$, and *edge set* $\bar{Q} - U = \bigcup_{e \in \mathcal{F}} \{ \gamma \text{ geometrically dual to } e \}$, where \mathcal{F} is the frontier of T in $V^{(1)}$ and “ $e \in \mathcal{F}$ ” means that e has an endpoint v such that $(e, v) \in \mathcal{F}$.

4. THE CENTERED DUAL VERSUS A DISK PACKING

We show in this section that the centered dual decomposition determined by \mathcal{S} interacts well with a set of disjoint open hyperbolic disks of equal radius isometrically embedded about the points of \mathcal{S} . Recall from the beginning of [3, §5] that a polygon P determines a *sector*




FIGURE 4.1. When a triangle determines a sector that it does not entirely contain.

of a disk U centered at one of its vertices x , with angle measure equal to $\angle_x P$, and that this sector contains $P \cap U$. Figure 4.1 illustrates an instance in which containment is proper, with the “bad” region shaded.

The *radius- R defect* $D_R(d_0, \dots, d_{n-1})$, as defined in [3, Definition 5.1], describes the area of the region of a cyclic n -gon P represented by $(d_0, \dots, d_{n-1}) \in \widetilde{\mathcal{AC}}_n$ outside the union of a collection of disjoint radius- R disks centered at its vertices, if each disk intersects P in a full sector. If P is centered then by [3, Lemma 5.3], the full sectors hypothesis holds, and by [3, Lemma 5.4] the area in question is $D_R(P)$. For non-centered cyclic polygons the pathology of Figure 4.1 may occur, but we show here that it does not for centered dual 2-cells.

Definition 4.1. For $\mathcal{S} \subset F$ finite, where F is a closed hyperbolic surface, define the *injectivity radius* $i(\mathcal{S})$ of F at \mathcal{S} to be the injectivity radius of the preimage $\widetilde{\mathcal{S}} \subset \mathbb{H}^2$, as defined above Lemma 1.1, of \mathcal{S} under the universal cover $\mathbb{H}^2 \rightarrow F$.

It is easy to see that $i(\mathcal{S})$ is the maximal R such that a collection of open, radius- R hyperbolic disks may be isometrically embedded in F without overlapping, centered at the points of \mathcal{S} . In particular, if $\mathcal{S} = \{x\}$ is a singleton, then $i(\mathcal{S})$ is the usual injectivity radius of F at x .

Proposition 4.2. Let V be the Voronoi tessellation and P_c the centered dual decomposition determined by $\mathcal{S} \subset F$ finite, where F is a closed hyperbolic surface. If $\{U_x\}$ is a set of open hyperbolic disks of radius $R \leq i(\mathcal{S})$ centered at the points of \mathcal{S} , then for a 2-cell Q of P_c :

$$\text{area} \left(Q - \bigcup_{x \in \mathcal{S}} (U_x \cap Q) \right) = D_R(Q) \doteq \sum_{v \in Q \cap V^{(0)}} D_R(\mathbf{d}_v),$$

where \mathbf{d}_v represents P_v in $\widetilde{\mathcal{AC}}_{n_v}$ for each v (with n_v the valence of v in $V^{(1)}$).

As we remarked above, this does not necessarily hold for Delaunay 2-cells that are non-centered; however, for those that are it follows directly from [3, Lemma 5.4]. Lemma 3.12 implies that each 2-cell of P_c is either a centered Delaunay polygon or contains a component T of $V_n^{(1)}$, so it is this latter case that we will address in the remainder of the section.

A centered dual 2-cell is by definition equal to the union of vertex polygons P_v for $v \in Q \cap V^{(0)}$. It will be convenient for our purposes to re-tile Q by a new set of “polygons.”

Definition 4.3. Let \tilde{V} be the Voronoi tessellation determined by $\tilde{\mathcal{S}} \subset \mathbb{H}^2$ closed and discrete, and let T be a component of $\tilde{V}_n^{(1)}$. For $v \in T^{(0)} - \{v_T\}$, let e_v be the edge of T with initial vertex v . For $v \in T^{(0)}$, define $v + 1$ to be the set of $w \in T^{(0)}$ such that v is the terminal vertex of e_w (oriented as in Definition 3.1).

Define $P'_{v_T} = \overline{P_v - (\bigcup_{w \in v+1} T(e_w, w))}$, where $T(e_w, w)$ is as in Definition 3.13, and if $v \neq v_T$ let $P'_v = \overline{(P_v \cup T(e_v, v)) - (\bigcup_{w \in v+1} T(e_w, w))}$ (here the overline denotes the closure in \mathbb{H}^2).

Although P'_v is not necessarily convex, its angle at a vertex x of P_v is clearly at most that of $P_v \cup T(e_v, v)$, so since $P_v \cup T(e_v, v)$ is convex (cf. Lemma 3.3) it makes sense to talk about “the sector determined by P'_v ” of a disk U centered at x . The key advantage of the P'_v is that they behave well with respect to such disks.

Lemma 4.4. *Let V be the Voronoi tessellation and P_c the centered dual decomposition determined by $\tilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$, where $\pi: \mathbb{H}^2 \rightarrow F$ is the universal cover of a closed hyperbolic surface and $\mathcal{S} \subset F$ is finite. Fix a component T of $V_n^{(1)}$, $v \in T^{(0)}$, and $R \leq i(\mathcal{S})$. A disk U_x of radius R centered at a vertex x of P intersects P'_v in the sector determined by P'_v . For $w \in v + 1$, $U_x \cap T(e_w, w) \neq \emptyset$ if and only if x is in the geometric dual γ_w to e_w .*

Let us recall that for v as above and $w \in v + 1$, $T(e_w, w) \subset T(e_w, v)$ and $T(e_w, v) \cap \partial T(e_w, v)$ is the edge γ_w geometrically dual to e_w , by Lemma 3.14. Thus Lemma 3.15 implies that $T(e_w, w)$ is entirely contained in $P_v \cup T(e_v, v)$, for $v \neq v_T$, or in P_v if $v = v_T$; and furthermore that $T(e_w, w) \cap \partial(P_v \cup T(e_v, v)) = \gamma_w$ (or that $T(e_w, w) \cap \partial P_v = \gamma_w$ if $v = v_T$).

Proof. For now take $v = v_T$. Lemma 3.15 implies that a vertex x of P is contained in $T(e, v)$, for some edge e containing v , if and only if x is an endpoint of the edge γ of P_v that is geometrically dual to e . Thus a small-enough disk U around x has the property that $U \cap P_v = (U \cap T(e, v)) \cup (U \cap T(e', v))$, where e and e' are the edges containing v with geometric duals γ and γ' meeting at x (this also uses Lemma 3.15).

For U_x as described above, [3, Lemma 5.2] implies that $U_x \cap T(e, v)$ is the sector determined by $T(e, v)$, and likewise for $T(e', v)$. Since $U_x \cap P_v$ is contained in the sector determined by P_v , and this is the union of those determined by $T(e, v)$ and $T(e', v)$ by the above, it follows that $U_x \cap P_v = (U_x \cap T(e, v)) \cup (U_x \cap T(e', v))$. For $w \in v + 1$ such that $e_w \neq e, e'$, since $T(e_w, w) \subset T(e_w, v)$ it follows that $T(e_w, w) \cap U_x = \emptyset$. On the other hand, if $e_w = e$ (say), then $\gamma_w = \gamma$ and U_x clearly intersects $T(e_w, w)$, in the sector that it determines (by [3, Lemma 5.2] again). The final assertion of the lemma follows.

Since $v = v_T$, the definition of P'_v implies that $U_x \cap P'_v = \overline{(U_x \cap P_v) - (\bigcup_{w \in v+1} (U_x \cap T(e_w, w)))}$. By the above, $U_x \cap P_v$ is a sector, and for each $w \in v + 1$, $U_x \cap T(e_w, w)$ is empty unless $x \in \gamma_w$, in which case it is a sector containing the boundary edge $U_x \cap \gamma_w$ of $U_x \cap P_v$. It thus easily follows from the description above that $U_x \cap P'_v$ is also a sector.

We have proved the lemma for $v = v_T$. The proof for $v \in T^{(0)} - \{v_T\}$ is similar, but with two important differences. First, P_v should be replaced above by $P_v \cup T(e_v, v)$, and second, the case of $x \in \partial \gamma_v$ must be treated separately. For such x it turns out that

$U_x \cap (P_v \cup T(e_v, v)) = U_x \cap T(e', v)$, where e' is the geometric dual to the edge of P_v meeting γ_v at x . This follows from [3, Lemma 5.3], which implies that $\angle_x(P_v \cup T(e_v, v)) = \angle_x T(e', v)$ after chasing through some definitions. \square

Fact. For each $v \in T^{(0)}$ and $w \in v + 1$, $P'_v \cap T(e_w, w)$ is the union of edges of $T(e_w, w)$ containing w . For $v \in T^{(0)} - \{v_T\}$, if γ_v is the geometric dual to e_v then:

$$\text{int } P'_v = (\text{int } P_v \cup \text{int } T(e_v, v) \cup \text{int } \gamma_v) - \left(\bigcup_{w \in v_T + 1} T(e_w, w) \right)$$

Similarly, $\text{int } P'_{v_T} = \text{int } P_{v_T} - (\bigcup_{w \in v_T + 1} T(e_w, w))$.

Lemma 4.5. *Let V be the Voronoi tessellation and P_c the centered dual decomposition determined by $\tilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$, where $\pi: \mathbb{H}^2 \rightarrow F$ is the universal cover of a closed hyperbolic surface and $\mathcal{S} \subset F$ is finite. If Q is a 2-cell of P_c containing a component T of $V_n^{(1)}$ then $Q = \bigcup_{v \in T^{(0)}} P'_v$, for P'_v as in Definition 4.3. For distinct v and w in $T^{(0)}$, $\text{int } P'_v \cap P'_w = \emptyset$.*

Proof. For $x \in Q$ let v_0 be such that $x \in P_{v_0}$. If $x \notin T(e_v, v)$ for any $v \in v_0 + 1$, then $x \in P'_{v_0}$. If there exists $v_1 \in v_0 + 1$ such that $x \in T(e_{v_1}, v_1)$, we choose $v \in T^{(0)}$ to satisfy three criteria: $x \in T(e_v, v)$, the reduced edge path from v_0 to v contains v_1 , and this edge path is longest among all of those joining v_0 to vertices satisfying the first two criteria. Since v_1 satisfies the first two criteria there is some such v , and by construction $x \in T(e_v, v)$ but not in $T(e_w, w)$ for any $w \in v + 1$ (the reduced edge path from v_0 to such a w is the union of the reduced edge path to v with e_w). Therefore $x \in P'_v$.

Our goal in the remainder is to prove that $\text{int } P'_v \cap P'_w = \emptyset$ for distinct v and w in $T^{(0)}$. For $v \in T^{(0)} - \{v_T\}$, let $e_0 \cup \dots \cup e_{n-1}$ be the reduced edge path in T joining v_T to v , numbered so that v_T is a vertex of e_0 and v is a vertex of e_{n-1} . Upon orienting this path and its edges as described in Lemma 3.7, v_T is the terminal vertex of e_0 and v the initial vertex of e_{n-1} . For $0 \leq i \leq n-1$ let v_i be the terminal vertex of e_i , so in particular $v_T = v_0$, and let $v_n = v$. Then v_i is the initial vertex of e_{i-1} for $i > 0$, so for each $i < n$, $e_i = e_{v_{i+1}}$ as defined in 4.3.

Claim 4.5.1. For $0 < i \leq n$, $T(e_{i-1}, v_i) \subset \bigcup_{j=0}^{i-1} P_{v_j}$; and if γ_j is the geometric dual to e_j for each j , we have $T(e_{i-1}, v_i) \cap \partial P_{v_0} \subset \gamma_0$, and $T(e_{i-1}, v_i) \cap \partial P_{v_j} \subset \gamma_{j-1} \cup \gamma_j$ for $0 < j \leq i-1$.

First take $i = 1$. Lemma 3.14 implies that $T(e_0, v_1) \subset T(e_0, v_0)$, and Lemma 3.15 implies that $T(e_0, v_0) \subset P_{v_0}$. (Recall from Lemma 3.7 that $P_{v_T} = P_{v_0}$ is centered.) By Definition 3.13, γ_0 is the edge of $T(e_0, v_0)$ opposite v_0 . Any other edge γ' of P_{v_0} is contained in $T(e', v_0)$, where e' is the geometric dual to γ' , and Lemma 3.15 implies that $T(e_0, v_0) \cap \gamma' \subset T(e_0, v_0) \cap T(e', v_0)$ is at most the endpoint $\gamma_0 \cap \gamma'$. Thus $T(e_0, v_0) \cap \partial P_{v_0} \subset \gamma_0$.

For $1 < i \leq n$, the combination of Lemmas 3.14 and 3.15 implies:

$$(4.5.2) \quad T(e_{i-1}, v_i) \subset T(e_{i-1}, v_{i-1}) \subset P_{v_{i-1}} \cup T(e_{i-2}, v_{i-1})$$

We assume by induction that $T(e_{i-2}, v_{i-1}) \subset \bigcup_{j=0}^{i-2} P_{v_j}$, so by the above $T(e_{i-1}, v_i) \subset \bigcup_{j=0}^{i-1} P_{v_j}$.

To prove the claim it remains to show that $T(e_{i-1}, v_i)$ has reasonable intersections with the P_{v_j} . We assume by induction that $T(e_{i-2}, v_{i-1}) \cap \partial P_{v_0} \subset \gamma_0$ and $T(e_{i-2}, v_{i-1}) \cap \partial P_{v_j} \subset \gamma_{j-1} \cup \gamma_j$ for $0 < j \leq i-2$. Thus by (4.5.2) it suffices to show that $T(e_{i-1}, v_{i-1}) \cap \partial P_{v_{i-1}} \subset \gamma_{i-2} \cup \gamma_{i-1}$. This follows as in the base case, but using the non-centered case of Lemma 3.15.

Below we will obtain different information from essentially the same sequence of observations.

Claim 4.5.3. For $n > 1$ and $0 < j < i \leq n$, $T(e_{i-1}, v_i) \subset T(e_{j-1}, v_j) \cup \left(\bigcup_{k=j}^{i-1} P_{v_k} \right)$.

We again prove the claim by induction, this time on n . For $n = 2$ the only case above is with $i = n = 2$ and $j = 1$. Applying (4.5.2) immediately implies the conclusion in this case. Let us now take $n > 2$ and suppose that the claim holds for $n - 1$. The only new cases to consider have $i = n$, since for $i \leq n - 1$ the conclusion follows from the induction hypothesis.

Fixing $i = n$, the conclusion for $j = n - 1$ is a direct application of (4.5.2). For $j < n - 1$, (4.5.2) gives $T(e_{n-1}, v_n) \subset T(e_{n-2}, v_{n-1}) \cup P_{v_{n-1}}$, so induction produces:

$$T(e_{n-1}, v_n) \subset T(e_{j-1}, v_j) \cup \left(\bigcup_{k=j}^{n-2} P_{v_k} \right) \cup P_{v_{n-1}} = T(e_{j-1}, v_j) \cup \left(\bigcup_{k=j}^{n-1} P_{v_k} \right)$$

This proves the claim.

We first fix $w = v_T$ and $v \in T^{(0)} - \{v_T\}$, and prove $\text{int } P'_v \cap P_w = P'_v \cap \text{int } P_w = \emptyset$. Let $e_0 \cup \dots \cup e_{n-1}$ be the reduced edge path joining v_T to v in T , numbered and oriented as above, so that $w = v_T$ is the terminal vertex of e_0 and v is the initial vertex of e_{n-1} . With the v_j also numbered as above for $0 \leq j \leq n$, we apply Claim 4.5.3 with $i = n$ and $j = 1$. Taking a union with P_v on each side of the result yields:

$$(4.5.4) \quad P_v \cup T(e_{n-1}, v) \subset T(e_0, v_1) \cup \left(\bigcup_{k=1}^n P_{v_k} \right)$$

This is because $v = v_n$. Since $v_1 \in v_T + 1$ and $e_0 = e_{v_1}$, the fact above the lemma implies that $\text{int } P'_{v_T} \cap T(e_0, v_1) = \emptyset$. By Lemma 1.9, $\text{int } P_{v_T} \cap P_{v_k} = \emptyset$ for $1 \leq k \leq n$. Since $P'_v \subset P_v \cup T(e_{n-1}, v)$, it follows from (4.5.4) that $P'_v \cap \text{int } P'_{v_T} = \emptyset$.

Since $P'_v \subset P_v \cup T(e_{n-1}, v)$, $\text{int } P'_v \subset \text{int } (P_v \cup T(e_{n-1}, v)) = \text{int } P_v \cup \text{int } T(e_{n-1}, v) \cup \text{int } \gamma_{n-1}$. The latter equality here follows from Lemma 3.15, which asserts that $T(e_{n-1}, v)$ intersects P_v in precisely the edge γ_{n-1} geometrically dual to e_{n-1} . Similarly:

Claim 4.5.5. For $0 < i \leq n$,

$$\text{int } (P_v \cup T(e_{n-1}, v_n)) \subset \text{int } T(e_{i-1}, v_i) \cup \left(\bigcup_{k=i}^n \text{int } P_{v_k} \right) \cup \left(\bigcup_{k=i-1}^{n-1} \text{int } \gamma_k \right),$$

where γ_k is the geometric dual to e_k for each k .

This uses Claim 4.5.3, also noting that $T(e_{i-1}, v) \cap \partial P_{v_k} \subset \gamma_{k-1} \cup \gamma_k$ for $i \leq k < n$, by Claim 4.5.1. Thus by the fact above the lemma, $\text{int } P'_v \cap P'_{v_T} = \emptyset$.

For distinct v and w in $T^{(0)} - \{v_T\}$, let $e_0 \cup \dots \cup e_{n-1}$ and $f_0 \cup \dots \cup f_{m-1}$ be reduced edge paths joining v_T to v and w , respectively. Let the e_i be numbered and oriented as in the case $w = v_T$ above, and number the vertices v_i accordingly. Similarly, number the f_i so that v_T is an endpoint of f_0 and w an endpoint of f_{m-1} , and orient them as in Lemma 3.7. Let w_i be the terminal vertex of f_i for each i , so $v_T = w_0$ in particular, and take $w = w_m$.

Because T is a tree, there are three cases to consider: either these paths have no edges in common, meeting only at v_T ; or one is an initial segment of the other; or they share an initial segment that is proper in each. We will address the third case in detail. The others are similar, and we will indicate afterwards how to approach them.

If $e_0 \cup \dots \cup e_{n-1}$ and $f_0 \cup \dots \cup f_{m-1}$ share an initial segment that is proper in each, then m and n are each at least 2. Let $i_0 > 0$ but less than $\min\{m, n\}$ be such that $e_i = f_i$ for $i < i_0$ but $e_{i_0} \neq f_{i_0}$. It follows that $v_i = w_i$ for $i \leq i_0$, but that $\{v_{i_0+1}, \dots, v_n\}$ is disjoint from $\{w_{i_0+1}, \dots, w_m\}$, since T is a tree.

Applying Claim 4.5.3 to $e_0 \cup \dots \cup e_{n-1}$ with $i = n$ and $j = i_0 + 1$, then taking the union with $P_v = P_{v_n}$ on both sides, we have:

$$(4.5.6) \quad P_v \cup T(e_{n-1}, v) \subset T(e_{i_0}, v_{i_0+1}) \cup \left(\bigcup_{k=i_0+1}^n P_{v_k} \right)$$

Using Claim 4.5.3 on $f_0 \cup \dots \cup f_{m-1}$ with $i = m$ and $j = i_0 + 1$, arguing as above yields:

$$(4.5.7) \quad P_w \cup T(f_{m-1}, w) \subset T(f_{i_0}, w_{i_0+1}) \cup \left(\bigcup_{k=i_0+1}^m P_{w_k} \right)$$

We will use Claim 4.5.5 with $i = i_0 + 1$ to show that $\text{int}(P_v \cup T(e_{n-1}, v))$ is disjoint from $P_w \cup T(f_{m-1}, w)$, from which it immediately follows that $\text{int } P'_v \cap P'_w = \emptyset$.

Lemma 1.9 implies that $\text{int } P_{v_k} \cap P_{w_l} = \emptyset$ for each $k \in \{i_0 + 1, \dots, n\}$ and $l \in \{i_0 + 1, \dots, m\}$, and also that $\text{int } \gamma_k \cap P_{w_l} = \emptyset$ for $k \in \{i_0, \dots, n-1\}$ and the same l . This is because $\gamma_k = P_{v_k} \cap P_{v_{k+1}}$ for each such k , and so its interior is disjoint from all P_v but these two. In the particular case $k = i_0$, γ_{i_0} is a different edge of $P_{v_{i_0}}$ than its edge of intersection with $P_{w_{i_0+1}}$, since $w_{i_0+1} \neq v_{i_0+1}$, so it is still true that $\text{int } \gamma_{i_0} \cap P_{w_{i_0+1}} = \emptyset$.

Lemma 3.14 implies that $T(e_{i_0}, v_{i_0+1}) \subset T(e_{i_0}, v_{i_0})$ and $T(f_{i_0}, w_{i_0+1}) \subset T(f_{i_0}, v_{i_0})$ (recall that $w_{i_0} = v_{i_0}$), and by Lemma 3.15 these are each contained in $P_{v_{i_0}} \cup T(e_{i_0-1}, v_{i_0})$. It further implies that $T(e_{i_0}, v_{i_0}) \cap T(f_{i_0}, v_{i_0})$ is at most an edge of each containing v_{i_0} ; hence in particular $\text{int } T(e_{i_0}, v_{i_0+1}) \cap T(f_{i_0}, w_{i_0+1}) = \emptyset$.

We will finally show that $\text{int } T(e_{i_0}, v_{i_0+1}) \cap P_{w_l} = \emptyset$ for each $l \in \{i_0 + 1, \dots, m\}$. Claim 4.5.1 implies that $T(e_{i_0}, v_{i_0+1}) \subset \bigcup_{j=0}^{i_0} P_{v_j}$, and the second part of that claim implies that its interior is contained in $\left(\bigcup_{j=0}^{i_0} \text{int } P_{v_j} \right) \cup \left(\bigcup_{j=0}^{i_0-1} \text{int } \gamma_j \right)$. It thus follows as above that $\text{int } T(e_{i_0}, v_{i_0+1}) \cap P_{w_l} = \emptyset$ for each l under consideration. This completes the proof that $\text{int } P'_v \cap P'_w = \emptyset$ when the edge paths joining each to v_T share a proper initial segment.

The case when the edge paths meet at v_T and nowhere else is similar to the above but with $i_0 = 0$. Then $T(e_0, v_1)$ and $T(f_0, w_1)$ are each contained in P_{v_T} , and it is immediate that

$\text{int } T(e_0, v_1)$ is disjoint from P_{w_l} for each $l > 0$. The case when $e_0 \cup \dots \cup e_{n-1}$ (say) is a proper initial segment of $f_0 \cup \dots \cup f_{m-1}$ is similar to the case $v = v_T$ that we first addressed. An extra argument is required in this case to show that $T(e_v, v) = T(e_{n-1}, v_n)$ has interior disjoint from the P_{w_l} ; this proceeds as in the paragraph directly above. \square

Definition 4.6. Let V be the Voronoi tessellation and P_c the centered dual decomposition determined by $\tilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$, where $\pi: \mathbb{H}^2 \rightarrow F$ is the universal cover of a closed hyperbolic surface and $\mathcal{S} \subset F$ is finite, and fix a 2-cell Q of P_c containing a component T of $V_n^{(1)}$. Define $T_0^{(0)} = \{v_T\}$ and $Q_0 = P_{v_T}$, and for $k > 0$ let $T_k^{(0)}$ consist of vertices of T joined to v_T by a path of at most k edges, and $Q'_k = \bigcup_{v \in T_k^{(0)}} P'_v$.

If x is a vertex of Q , for $k \geq 0$ define the *restriction* of U_x to Q'_k as the union of sectors:

$$U_x|_{Q'_k} \doteq \bigcup \{U_x \cap P'_v \mid v \in T_k^{(0)}, x \in P_v\}$$

Let the restriction $U_x|_Q$ of U_x to Q be the union of its restrictions to Q'_k over all $k \geq 0$.

The point of defining the restriction is to exclude an incidental component of intersection with Q'_k as in Figure 4.1, where a disk U_x protrudes from a polygon that does not entirely contain the sector that it determines, intersecting one that does not have x as a vertex.

Lemma 4.7. *Let V be the Voronoi tessellation and P_c the centered dual decomposition determined by $\tilde{\mathcal{S}} = \pi^{-1}(\mathcal{S})$, where $\pi: \mathbb{H}^2 \rightarrow F$ is the universal cover of a closed hyperbolic surface and $\mathcal{S} \subset F$ is finite. If $\{U_x\}$ is a set of open hyperbolic disks of radius $R \leq i(\mathcal{S})$ centered at the vertices of a 2-cell Q of P_c containing a component T of $V_n^{(1)}$, for each $k \geq 0$:*

$$\text{area} \left(Q'_k - \bigcup_{x \in \mathcal{S}} U_x|_{Q'_k} \right) = \sum_{v \in T_k^{(0)}} D_R(\mathbf{d}_v) - \sum_{w \in T_{k+1}^{(0)} - T_k^{(0)}} D_R(d_{e_w}, J_w),$$

where \mathbf{d}_v represents P_v in $\tilde{\mathcal{AC}}_{n_v}$ for each v (with n_v the valence of v in $V^{(1)}$), and $D_R(d_{e_w}, J_w)$ is as in [3, Lemma 5.2] (with d_{e_w} the length of the geometric dual to e_w).

Proof. We prove this by induction on k . By definition, $P_{v_T} = P'_{v_T} \cup (\bigcup_{w \in v_T+1} T(e_w, w))$. For each vertex x of P_{v_T} , $U_x \cap P_{v_T}$ is a sector (by [3, Lemma 5.3]) that is a non-overlapping union of sectors $(U_x \cap P'_{v_T}) \cup (U_x \cap T(e_w, w))$ (by Lemma 4.4 and [3, Lemma 5.2], respectively). By Lemma 4.4, $T(e_w, w)$ does not intersect U_x unless x is one of its vertices. Thus:

$$\text{area} \left(P'_{v_T} - \bigcup_{x \in \mathcal{S}} U_x|_{Q'_0} \right) = D_R(\mathbf{d}_{v_T}) - \sum_{w \in v_T+1} D_R(d_{e_w}, J_w),$$

This is because [3, Lemma 5.4] implies that $D_R(\mathbf{d}_{v_T})$ is the area of the region in P_{v_T} outside the U_x , and [3, Lemma 5.2] implies that $D_R(d_{e_w}, J_w)$ is the area of the region of $T(e_w, w)$ outside the disks centered at its endpoint. Since $T^{(0)} = \{v_T\}$ and $v_T + 1 = T_1^{(0)} - T_0^{(0)}$, the $k = 0$ case follows.

For $k > 0$, we note that for any $v \in T_k^{(0)}$, $\text{int } P'_v$ intersects $U_x|_{Q'_k}$ if and only if x is a vertex of P_v . This by definition of the $U_x|_{Q'_k}$, since by Lemma 4.4 for a vertex x' of P'_w for some $w \neq v$, P'_w contains the sector of $U_{x'}$ that it determines, and $\text{int } P'_v \cap P'_w = \emptyset$ by Lemma 4.5.

For $k > 0$, assume the conclusion holds for $k - 1$. Writing $Q'_k = Q'_{k-1} \cup \left(\bigcup_{v \in T_k^{(0)} - T_{k-1}^{(0)}} P'_v \right)$ yields:

$$\begin{aligned} \text{area} \left(Q'_k - \bigcup_{x \in \mathcal{S}} U_x|_{Q'_k} \right) &= \sum_{v \in T_{k-1}^{(0)}} D_R(\mathbf{d}_v) - \sum_{w \in T_k^{(0)} - T_{k-1}^{(0)}} D_R(d_{e_w}, J_w), \\ &\quad + \sum_{v \in T_k^{(0)} - T_{k-1}^{(0)}} \left(D_R(\mathbf{d}_v) + D_R(d_{e_v}, J_v) - \sum_{w \in v+1} D_R(d_{e_w}, J_w) \right) \end{aligned}$$

The first line above follows from the inductive hypothesis, and the second by an argument analogous to the base case. The sum above telescopes, and since $T_{k+1}^{(0)} - T_k^{(0)} = \bigcup \{v + 1 \mid v \in T_k^{(0)} - T_{k-1}^{(0)}\}$, the lemma follows by induction. \square

The main result of the section follows quickly.

Proof of Proposition 4.2. The result follows from Lemma 4.7 and two observations.

First, for a 2-cell Q of P_c containing a component T of $V_n^{(1)}$, let \tilde{Q} be a lift to \mathbb{H}^2 and \tilde{T} the lift of T that it contains. Because \tilde{T} is finite, there exists $k > 0$ such that $\tilde{T} = \tilde{T}_k^{(0)}$, and hence $\tilde{Q} = \tilde{Q}'_k$ (by Lemma 4.5) and $\tilde{T}_{k+1}^{(0)} - \tilde{T}_k^{(0)} = \emptyset$. Thus Lemma 4.7 implies:

$$\text{area} \left(\tilde{Q} - \bigcup_{x \in \tilde{\mathcal{S}}} U_x|_{\tilde{Q}} \right) = \sum_{v \in \tilde{T}^{(0)}} D_R(\mathbf{d}_v)$$

The first observation above, combined with Lemma 4.4, implies for each 2-cell \tilde{Q} of \tilde{P}_c that contains a component \tilde{T} of \tilde{V}_n^1 , and each $x \in \tilde{\mathcal{S}}$ that is the vertex of P_v for some $v \in \tilde{T}^{(0)}$, that \tilde{Q} contains the union of sectors of U_x determined by the P_v over all $v \in T^{(0)}$. The same holds for 2-cells \tilde{Q} that are centered polygons, by [3, Lemma 5.3]. It follows that for each $x \in \tilde{\mathcal{S}}$, $U_x \subset \bigcup_{i=1}^n Q_i$, where $\{Q_i\}_{i=1}^n$ is the collection of 2-cells of \tilde{P}_c containing x . This in turn implies the second observation: that $U_x \cap \tilde{Q} = U_x|_{\tilde{Q}}$ for each $x \in \tilde{\mathcal{S}}$. \square

5. ADMISSIBLE SPACES

This section is devoted to abstracting the data provided by a 2-cell Q of the centered dual and lower bounds for its edge lengths, turning this into a parameter space and a function on it whose minimum bounds the defect below. We will show that this defect function attains a minimum on the closure of the parameter space, and in the second half of the section restrict the location of such a minimum for low-complexity cells.

Definition 5.1. Let $T \subset V$ be finite graphs such that T is a rooted tree with root vertex v_T . Partially order $T^{(0)}$ by setting $v < v_T$ for each $v \in T^{(0)} - \{v_T\}$ and $w < v$ if the edge arc in T joining $w \in T^{(0)} - \{v_T, v\}$ to v_T runs through v . A vertex v is *minimal* if there is no $w \in T^{(0)}$ such that $w < v$; ie, so that $v+1$ (as in Definition 4.3) is empty. For $v \in T^{(0)} - \{v_T\}$, say “ $e \rightarrow v$ ” for each edge $e \neq e_v$ of V containing v , where e_v is as in Definition 4.3.

Definition 5.2. Let $T \subset V$ be finite graphs such that T is a rooted tree with root vertex v_T and each vertex of V has valence at least three. Let \mathcal{E} be the set of edges of T and \mathcal{F} the frontier of T in V , fix an ordering on $\mathcal{E} \cup \mathcal{F}$ and for some choice of $d_e > 0$ for each $e \in \mathcal{F}$, let $\mathbf{d}_{\mathcal{F}} = (d_e) \in \mathbb{R}^{\mathcal{F}}$. For any $\mathbf{d}_{\mathcal{E}} = (d_e) \in \mathbb{R}^{\mathcal{E}}$, let $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ and $P_v(\mathbf{d}) = (d_{e_0}, \dots, d_{e_{n_v-1}})$ for $v \in T^{(0)}$, where the edges of V containing v are cyclically ordered as e_0, \dots, e_{n_v-1} . We will say $\mathbf{d}_{\mathcal{E}}$ is in the *admissible set* $Ad(\mathbf{d}_{\mathcal{F}})$ determined by $\mathbf{d}_{\mathcal{F}}$ if the following criteria hold:

- (1) For $v \in T^{(0)} - \{v_T\}$ with valence n_v in V , $P_v(\mathbf{d}) \in \mathcal{AC}_{n_v} - \mathcal{C}_{n_v}$ has largest entry d_{e_v} .
- (2) $P_{v_T}(\mathbf{d}) \in \mathcal{C}_{n_T}$, where v_T has valence n_T in V .
- (3) $J(P_v(\mathbf{d})) > J(P_w(\mathbf{d}))$ for each $w \in v - 1$, where $J(P_v(\mathbf{d}))$ and $J(P_w(\mathbf{d}))$ are the respective radii of $P_v(\mathbf{d})$ and $P_w(\mathbf{d})$.

(Note that the final condition above is vacuous for minimal v .)

Definition 5.3. Let $T \subset V$ be finite graphs such that T is a rooted tree with root vertex v_T and each vertex of V has valence at least three. Let \mathcal{E} be the set of edges of T and \mathcal{F} the frontier of T in V , and fix $\mathbf{d}_{\mathcal{F}} = (d_e | e \in \mathcal{F}) \in \mathbb{R}^{\mathcal{F}}$ such that $Ad(\mathbf{d}_{\mathcal{F}}) \neq \emptyset$. For each $\mathbf{d}_{\mathcal{E}} \in Ad(\mathbf{d}_{\mathcal{F}})$, let $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ and for $R \leq \min\{d_e/2 | e \in \mathcal{F}\}$, define:

$$D_R(T, \mathbf{d}) = \sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d})),$$

where $P_v(\mathbf{d})$ is as in Definition 5.2 and $D_R(P)$ is as defined in [3, Lemma 5.4].

Lemma 5.4. *Let V be the Voronoi decomposition of a closed hyperbolic surface F determined by a finite set $\mathcal{S} \subset F$, and let Q be a centered dual 2-cell containing a component T of $V_n^{(1)}$. Let \mathcal{E} be the edge set of T and \mathcal{F} its frontier in V , and for each $e \in \mathcal{E}$ or such that $(e, v) \in \mathcal{F}$ for some v , let d_e be the length of the geometric dual to e . Then $\mathbf{d}_{\mathcal{E}} \in Ad(\mathbf{d}_{\mathcal{F}})$ and $D_R(Q) = D_R(T, \mathbf{d})$ for $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ and $R \leq \min\{d_e/2 | e \in \mathcal{F}\}$.*

Proof. For each $v \in T^{(0)}$, Lemma 1.9 implies that the vertex polygon P_v is a cyclic polygon with cyclically ordered side length collection $(d_{e_0}, \dots, d_{e_{n_v-1}})$, where e_0, \dots, e_{n_v-1} is the cyclically ordered collection of edges of V containing v . Recall from Definition 3.1 that each edge of T is oriented. Lemma 3.6 asserts that the root vertex v_T of T is the terminal endpoint of every edge of T that contains it. Since every other edge of V containing v_T is centered, Lemma 3.3 implies that $P_{v_T} \in \mathcal{C}_n$. This establishes (2) from Definition 5.2.

Lemma 3.3 also implies that for $v \in T^{(0)} - \{v_T\}$, P_v is non-centered with longest side length d_{e_v} , yielding (1) from Definition 5.2. For $v \in T^{(0)}$ and $w \in v - 1$, Lemma 3.7 implies that w is the initial vertex of e_w , and the definition (in 5.1 above) implies that v is its terminal vertex. Therefore Lemma 3.2 yields $J_v > J_w$ and hence, by Lemma 1.9, that $J(P_v) > J(P_w)$. This establishes (3) from Definition 5.2.

That $D_R(Q) = D_R(T, \mathbf{d})$ is a direct consequence of Definition 5.3 and Proposition 4.2. \square

The lemma below implies in particular that the admissible set of $\mathbf{d}_{\mathcal{E}}$ is a bounded subset of $\mathbb{R}^{\mathcal{F}}$, so that it has compact closure.

Lemma 5.5. *Let $T \subset V$ be finite graphs such that T is a rooted tree with root vertex v_T and each vertex of V has valence at least three. Let \mathcal{E} be the set of edges of T and \mathcal{F} the frontier of T in V . There exist collections $\{b_e: (\mathbb{R}^+)^{\mathcal{F}} \rightarrow \mathbb{R}^+ \mid e \in \mathcal{E}\}$ and $\{h_e: (\mathbb{R}^+)^{\mathcal{F}} \rightarrow \mathbb{R}^+ \mid e \in \mathcal{E}\}$ characterized by the following properties:*

- $P_v(\mathbf{d}) \in \widetilde{\mathcal{BC}}_{n_v}$, with largest entry d_{e_v} , for each $v \in T^{(0)} - \{v_T\}$, where v has valence n_v in V and $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$, if and only if $d_e = b_e(\mathbf{d}_{\mathcal{F}})$ for each $e \in \mathcal{E}$.
- $P_v(\mathbf{d}) \in \widetilde{\mathcal{HC}}_{n_v}$, with largest entry d_{e_v} , for each $v \in T^{(0)} - \{v_T\}$ if and only if $d_e = h_e(\mathbf{d}_{\mathcal{F}})$ for each $e \in \mathcal{E}$.

For fixed $\mathbf{d}_{\mathcal{F}}$, if $\mathbf{d}_{\mathcal{E}} = (d_e \mid e \in \mathcal{E}) \in \text{Ad}(\mathbf{d}_{\mathcal{F}})$ then $b_e(\mathbf{d}_{\mathcal{F}}) \leq d_e < h_e(\mathbf{d}_{\mathcal{F}})$ for each $e \in \mathcal{E}$; and for $\mathbf{d}'_{\mathcal{F}} = (d'_e \mid e \in \mathcal{F})$ such that $d'_e \geq d_e$ for each $e \in \mathcal{F}$, $b_e(\mathbf{d}'_{\mathcal{F}}) \geq b_e(\mathbf{d}_{\mathcal{F}})$ for each $e \in \mathcal{E}$.

Proof. The proof is by induction, the key point being that for $v \in T^{(0)} - \{v_T\}$, $b_{e_v}(\mathbf{d}_{\mathcal{F}})$ is determined by the set of $b_{e_w}(\mathbf{d}_{\mathcal{F}})$ for $w < v$, and similarly for $h_{e_v}(\mathbf{d}_{\mathcal{F}})$. Fix $\mathbf{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$.

Suppose first that $v \in T^{(0)}$ is minimal, so each $e \rightarrow v$ is in \mathcal{F} . Cyclically enumerate the edges of V containing v as e_0, \dots, e_{n-1} so that $e_0 = e_v$, and for each $i > 0$ let $d_i = d_{e_i}$. Then [3, Lemma 3.4] implies that $b_{e_v}(\mathbf{d}_{\mathcal{F}}) \doteq b_0(d_1, \dots, d_{n-1})$ is the unique number such that $(b_{e_v}(\mathbf{d}_{\mathcal{F}}), d_1, \dots, d_{n-1}) \in \overline{U}_n \cap \widetilde{\mathcal{BC}}_n$, where U_n is as in [3, Lemma 3.2]. That is, $b_{e_v}(\mathbf{d}_{\mathcal{F}})$ is unique with the property that the tuple above is in $\widetilde{\mathcal{BC}}_n$ and has its largest entry first. [3, Lemma 3.10] implies the analogous conclusion for $h_{e_v}(\mathbf{d}_{\mathcal{F}}) \doteq h_0(d_1, \dots, d_{n-1})$ and $\widetilde{\mathcal{HC}}_n$.

Let us also note that by [3, Corollary 3.11], $(d_0, d_1, \dots, d_{n-1}) \in U_n$ is in $\widetilde{\mathcal{AC}}_n - \widetilde{\mathcal{C}}_n$ if and only if $b_{e_v}(\mathbf{d}_{\mathcal{F}}) \leq d_0 < h_{e_v}(\mathbf{d}_{\mathcal{F}})$. If $\mathbf{d}_{\mathcal{E}} \in \text{Ad}(\mathbf{d}_{\mathcal{F}})$ then for $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ Definition 5.2 (1) implies that $P_v(\mathbf{d}) = (d_{e_v}, d_1, \dots, d_{n-1}) \in U_n \cap (\widetilde{\mathcal{AC}}_n - \widetilde{\mathcal{C}}_n)$, so $b_{e_v}(\mathbf{d}_{\mathcal{F}}) \leq d_{e_v} < h_{e_v}(\mathbf{d}_{\mathcal{F}})$.

Now fix $v \in T^{(0)} - \{v_T\}$ non-minimal, and suppose that we $b_{e_w}(\mathbf{d}_{\mathcal{F}})$ and $h_{e_w}(\mathbf{d}_{\mathcal{F}})$ are defined for each $w < v$, satisfying the following inductive hypotheses:

- $P_w(\mathbf{d}) \in \widetilde{\mathcal{BC}}_{n_w}$, with largest entry d_{e_w} , for each $w < v$, where $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$, if and only if $d_{e_w} = b_{e_w}(\mathbf{d}_{\mathcal{F}})$ for each $w < v$.
- $P_w(\mathbf{d}) \in \widetilde{\mathcal{BC}}_{n_w}$, with largest entry d_{e_w} , for each $w < v$, where $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$, if and only if $d_{e_w} = b_{e_w}(\mathbf{d}_{\mathcal{F}})$ for each $w < v$.
- For each $\mathbf{d}_{\mathcal{E}} \in \text{Ad}(\mathbf{d}_{\mathcal{F}})$, $b_{e_w}(\mathbf{d}_{\mathcal{F}}) \leq d_{e_w} < h_{e_w}(\mathbf{d}_{\mathcal{F}})$ for each $w < v$.

Cyclically enumerate the edges containing v as e_0, \dots, e_{n-1} so that $e_0 = e_v$, and for $i > 0$ define:

$$d_i = \begin{cases} d_{e_i} & e_i \in \mathcal{F} \\ b_{e_i}(\mathbf{d}_{\mathcal{F}}) & e_i \in \mathcal{E} \end{cases} \quad d'_i = \begin{cases} d_{e_i} & e_i \in \mathcal{F} \\ h_{e_i}(\mathbf{d}_{\mathcal{F}}) & e_i \in \mathcal{E} \end{cases}$$

Then [3, Lemma 3.4] again implies that $b_{e_v}(\mathbf{d}_{\mathcal{F}}) \doteq b_0(d_1, \dots, d_{n-1})$ is the unique number such that $(b_{e_v}(\mathbf{d}_{\mathcal{F}}), d_1, \dots, d_{n-1}) \in \overline{U}_n \cap \widetilde{\mathcal{BC}}_n$, and [3, Lemma 3.10] gives the analogous conclusion for $h_{e_v}(\mathbf{d}_{\mathcal{F}}) = h_0(d'_1, \dots, d'_{n-1})$.

For $\mathbf{d}_{\mathcal{E}} \in \text{Ad}(\mathbf{d}_{\mathcal{F}})$, by hypothesis $d_i \leq d_{e_i}$ for each $i \in \{1, \dots, n-1\}$, so [3, Lemma 3.4] implies that $b_{e_v}(\mathbf{d}_{\mathcal{F}}) \leq b_0(d_{e_1}, \dots, d_{e_{n-1}})$. We also have $d_{e_i} < d_i$ for each i such that $e_i \in \mathcal{E}$ by hypothesis

(and $d_{e_i} = d_i$ otherwise), so $h_{e_v}(\mathbf{d}_{\mathcal{F}}) > h_0(d_{e_1}, \dots, d_{e_{n-1}})$ by [3, Lemma 3.10]. By Definition 5.2 (1) and [3, Corollary 3.11],

$$b_0(d_{e_1}, \dots, d_{e_{n-1}}) \leq d_{e_v} < h_0(d_{e_1}, \dots, d_{e_{n-1}}),$$

and it follows that $b_{e_v}(\mathbf{d}_{\mathcal{F}}) \leq d_{e_v} < h_{e_v}(\mathbf{d}_{\mathcal{F}})$. We have thus proved the three hypotheses above for $\{v\} \cup \{w < v\}$, so it follows by induction that they hold on all of $T^{(0)} - \{v_T\}$. (Recall in particular that there is a unique e_v for each $v \in T^{(0)} - \{v_T\}$, and that \mathcal{E} is the set of all such e_v .)

The final claim of the lemma, that b_e is “increasing” in $\mathbf{d}_{\mathcal{F}}$ for each e , follows from an inductive argument and [3, Lemma 3.4], which asserts that $b_0(d_1, \dots, d_{n-1}) \leq b_0(d'_1, \dots, d'_{n-1})$ when $d_i \leq d'_i$ for each i . \square

Remark 5.6. For any given tree T with frontier \mathcal{F} , the proof of Lemma 5.5 is easily adapted (using formulas from [3]) to produce a recursive algorithm that takes $\mathbf{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ and computes the values $b_e(\mathbf{d}_{\mathcal{F}})$ or $h_e(\mathbf{d}_{\mathcal{F}})$ from the “outside in.”

Lemma 5.7. *Let $T \subset V$ be finite graphs such that T is a rooted tree with root vertex v_T and each vertex of V has valence at least three. Let \mathcal{E} be the set of edges of T and \mathcal{F} the frontier of T in V , and fix $\mathbf{d}_{\mathcal{F}} = (d_e \mid e \in \mathcal{F}) \in \mathbb{R}^{\mathcal{F}}$. If $\text{Ad}(\mathbf{d}_{\mathcal{F}}) \neq \emptyset$, then for each $\mathbf{d}_{\mathcal{E}}$ in its closure $\overline{\text{Ad}}(\mathbf{d}_{\mathcal{F}})$, $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ satisfies:*

- (1) *For $v \in T^{(0)} - \{v_T\}$ with valence n_v in V , $P_v(\mathbf{d}) \in \widetilde{\mathcal{AC}}_{n_v} - \widetilde{\mathcal{C}}_{n_v}$ has largest entry d_{e_v} .*
- (2) *$P_{v_T}(\mathbf{d}) \in \widetilde{\mathcal{C}}_{n_T} \cup \widetilde{\mathcal{BC}}_{n_T}$, where v_T has valence n_T in V .*
- (3) *$J(P_v(\mathbf{d})) \geq J(P_w(\mathbf{d}))$ for each $w \in v - 1$, where $J(P_v(\mathbf{d}))$ and $J(P_w(\mathbf{d}))$ are the respective radii of $P_v(\mathbf{d})$ and $P_w(\mathbf{d})$.*

Proof. Lemma 5.5 implies that $\text{Ad}(\mathbf{d}_{\mathcal{F}})$ is bounded in $\mathbb{R}^{\mathcal{E}}$ and therefore has compact closure. Since $P_{v_T}(\mathbf{d}) \in \mathcal{C}_n$ for each $\mathbf{d} \in \text{Ad}(\mathbf{d}_{\mathcal{F}}) \times \{\mathbf{d}_{\mathcal{F}}\}$, for \mathbf{d} in the closure it must be the case that $P_{v_T}(\mathbf{d}) \in \overline{\mathcal{C}}_n$, establishing (2). [3, Proposition 4.1] implies that for each $v \in T^{(0)}$, $J(P_v(\mathbf{d}))$ varies continuously with \mathbf{d} on $\text{Ad}(\mathbf{d}_{\mathcal{F}}) \times \{\mathbf{d}_{\mathcal{F}}\}$. Thus since $J(P_v(\mathbf{d})) > J(P_w(\mathbf{d}))$ for each such \mathbf{d} and $w \in v - 1$, $J(P_v(\mathbf{d})) \geq J(P_w(\mathbf{d}))$ for each $\mathbf{d} \in \overline{\text{Ad}}(\mathbf{d}_{\mathcal{F}}) \times \{\mathbf{d}_{\mathcal{F}}\}$, so (3) holds.

Now suppose that (1) does not hold, so there exist $\mathbf{d} \in \overline{\text{Ad}}(\mathbf{d}_{\mathcal{F}}) \times \{\mathbf{d}_{\mathcal{F}}\}$ and $v \in T^{(0)} - \{v_T\}$ such that $P_v(\mathbf{d}) \in \widetilde{\mathcal{AC}}_{n_v} - \mathcal{AC}_{n_v}$. Let us take v to be maximal with this property, so that in particular $P_w(\mathbf{d}) \in \mathcal{AC}_n$ for the endpoint w of e_v . If $\{\mathbf{d}_n\}$ is a sequence in $\text{Ad}(\mathbf{d}_{\mathcal{F}}) \times \{\mathbf{d}_{\mathcal{F}}\}$ approaching \mathbf{d} , then $P_w(\mathbf{d}_n) \rightarrow P_w(\mathbf{d})$, and so there is a universal upper bound on $J(P_w(\mathbf{d}_n))$. On the other hand, [3, Lemma 4.7] implies that $J(P_v(\mathbf{d}_n)) \rightarrow \infty$, contradicting criterion (3) of Definition 5.2 for some \mathbf{d}_n . Therefore (1) holds. \square

Lemma 5.8. *Let $T \subset V$ be finite graphs such that T is a rooted tree with root vertex v_T and each vertex of V has valence at least three. Let \mathcal{E} be the set of edges of T and \mathcal{F} the frontier of T in V , and fix $\mathbf{d}_{\mathcal{F}} = (d_e \mid e \in \mathcal{F}) \in \mathbb{R}^{\mathcal{F}}$ such that $\text{Ad}(\mathbf{d}_{\mathcal{F}}) \neq \emptyset$. Then $D_R(T, \mathbf{d})$ is continuous on $\overline{\text{Ad}}(\mathbf{d}_{\mathcal{F}}) \times \{\mathbf{d}_{\mathcal{F}}\}$ and attains a minimum there.*

Proof. Since $P \mapsto D_R(P)$ is continuous on \mathcal{AC}_n (by [3, Proposition 5.5]), and by the above $P_v(\mathbf{d}) \in \mathcal{AC}_n$ for each $\mathbf{d} \in \overline{\text{Ad}}(\mathbf{d}_{\mathcal{F}}) \times \{\mathbf{d}_{\mathcal{F}}\}$, $D_R(T, \mathbf{d})$ is continuous on $\overline{\text{Ad}}(\mathbf{d}_{\mathcal{F}}) \times \{\mathbf{d}_{\mathcal{F}}\}$. This

set is closed and, by Lemma 5.5, bounded, so it is compact and $D_R(T, \mathbf{d})$ attains a minimum on it. \square

For an arbitrary finite tree T and $\mathbf{d}_{\mathcal{F}} \in \mathbb{R}^{\mathcal{F}}$ as above, it seems difficult to precisely describe $Ad(\mathbf{d}_{\mathcal{F}})$ or determine the point in $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$ at which $D_R(T, \mathbf{d})$ attains its minimum. Here we will identify an alternative that such a minimum point must satisfy, at least for very simple T : those with one or two edges. In the second half of the section, we will turn this into an algorithm that produces lower bounds on the minimum of $D_R(T, \mathbf{d})$, given lower bounds on the entries of $\mathbf{d}_{\mathcal{F}}$.

We first address the case that T has a single edge. In this case, uniquely, we are able to describe the topology of $Ad(\mathbf{d}_{\mathcal{F}})$ and locate the minimum of D_R .

Lemma 5.9. *Let V be a graph and T a subgraph with one edge e_T and root vertex v_T , and let \mathcal{F} be the frontier of T in V . For $\mathbf{d}_{\mathcal{F}} \in \mathbb{R}^{\mathcal{F}}$, if $Ad(\mathbf{d}_{\mathcal{F}}) \neq \emptyset$ it is an interval: (d^-, d^+) or $[d^-, d^+)$. For $R \geq 0$, $D_R(T, \mathbf{d})$ attains its minimum at $\mathbf{d} = (d^-, \mathbf{d}_{\mathcal{F}})$, which satisfies one of:*

- (1) $P_v(\mathbf{d}) \in \widetilde{\mathcal{BC}}_n$, where n is the valence in V of the initial vertex v of e_T ; or
- (2) $P_{v_T}(\mathbf{d}) \in \widetilde{\mathcal{BC}}_{n_T}$, where n_T is the valence of v_T in V .

In case (2) above, d^- is not the largest side length of $P_{v_T}(\mathbf{d})$.

Proof. By Definition 5.2 $Ad(\mathbf{d}_{\mathcal{F}})$ is contained in a subset of \mathbb{R}^+ consisting of possible values for d_{e_T} . By Lemma 5.5, if criterion (1) is satisfied then $d_{e_T} \in [b_{e_T}, h_{e_T})$, where $b_{e_T} = b_{e_T}(\mathbf{d}_{\mathcal{F}})$ and $h_{e_T} = h_{e_T}(\mathbf{d}_{\mathcal{F}})$ as in Lemma 5.5. In fact, $P_v(\mathbf{d}) \in \widetilde{\mathcal{AC}}_n - \widetilde{\mathcal{C}}_n$ if and only if $d_{e_T} \in [b_{e_T}, h_{e_T})$, where $\mathbf{d} = (d_{e_T}, \mathbf{d}_{\mathcal{F}})$. This follows from [3, Corollary 3.11], as pointed out in the base case of the proof of Lemma 5.5, and it follows that (1) is satisfied if and only if $d_{e_T} \in [b_{e_T}, h_{e_T})$.

Now consider criterion (2). Let the edges of V containing v_T be cyclically enumerated $e_0, e_1, \dots, e_{n_T-1}$ so that $e_0 = e_T$. Then $e_i \in \mathcal{F}$ for $1 \leq i < n_T$. Let $d = d_{e_T}$ and $d_i = d_{e_i}$ for $i > 0$, and for $\mathbf{d} = (d, \mathbf{d}_{\mathcal{F}})$ recall that from Definition 5.2 that:

$$P_{v_T}(\mathbf{d}) \doteq (d, d_1, \dots, d_{n_T-1}) \in (\mathbb{R}^+)^n$$

Let $M = \max\{d_i\}_{i=1}^{n_T-1}$. The inequality of [3, Definition 3.1], determining whether $P_{v_T}(\mathbf{d}) \in \widetilde{\mathcal{C}}_n$, takes different forms depending on the relation of d to M . For $d \geq M$, $P_{v_T}(\mathbf{d}) \in \mathcal{C}_{n_T}$ if and only if $A_d(d/2) + \sum_{i=1}^{n_T-1} A_{d_i}(d/2) > 2\pi$, where $A_d(J)$ is from [3, Lemma 1.7]. By [3, Lemma 1.8], $A_d(d/2) = \pi$ and $\sum_{i=1}^{n_T-1} A_{d_i}(J)$ decreases in J to a horizontal asymptote of 0, so the criterion of [3, Definition 3.1] is satisfied at $d = M$ and there exists $D_0^+ > M$ such that $A_d(d/2) + \sum_{i=1}^{n_T-1} A_{d_i}(d/2) > 2\pi$ if and only if $d < D_0^+$.

For $d \leq M$, [3, Definition 3.1] requires $A_d(M/2) + \sum_{i=1}^{n_T-1} A_{d_i}(M/2) > 2\pi$ for $P_v(\mathbf{d}) \in \widetilde{\mathcal{C}}_n$. Since $A_d(J)$ is continuous and increases in d there is an open interval of positive $d < M$, with left endpoint $D_0^- \geq 0$, on which this inequality holds. Thus $\{d \mid P_{v_T}(\mathbf{d}) \in \mathcal{C}_n\} = (D_0^-, D_0^+)$. If $D_0^- > 0$ then $A_{D_0^-}(M/2) + \sum_{i=1}^{n_T-1} A_{d_i}(M/2) = 2\pi$. In this case, for $\mathbf{d} = (D_0^-, \mathbf{d}_{\mathcal{F}})$ [3, Lemma 3.3] implies $P_{v_T}(\mathbf{d})$ is in the closure of $\widetilde{\mathcal{C}}_{n_T}$. It is not in $\widetilde{\mathcal{C}}_{n_T}$, so $P_{v_T}(\mathbf{d}) \in \widetilde{\mathcal{BC}}_{n_T}$ by [3, Lemma 3.4]. Furthermore, its longest side length is $M > D_0^-$, since $M = d_i$ for some $i > 0$.

If $Ad(\mathbf{d}_{\mathcal{F}}) \neq \emptyset$ then $[b_{e_T}, h_{e_T}) \cap (D_0^-, D_0^+)$ is non-empty. We claim that $f(d) = J(P_{v_T}(\mathbf{d})) - J(P_v(\mathbf{d}))$ decreases in d on $[b_{e_T}, h_{e_T}) \cap (D_0^-, D_0^+)$, where $J: \widetilde{\mathcal{AC}}_n \rightarrow \mathbb{R}^+$ is as in [3, Lemma 3.6]. This follows directly from [3, Lemma 4.5], which implies that on this interval:

$$\left| \frac{\partial}{\partial d} J(P_{v_T}(\mathbf{d})) \right| < \frac{1}{2} < \left| \frac{\partial}{\partial d} J(P_v(\mathbf{d})) \right|$$

By criterion (3) of Definition 5.2, $Ad(\mathbf{d}_{\mathcal{F}}) = f^{-1}(\mathbb{R}^+) \cap ([b_{e_T}, h_{e_T}) \cap (D_0^-, D_0^+))$, so since f is decreasing $Ad(\mathbf{d}_{\mathcal{F}})$ is a subinterval containing the left endpoint of $[b_{e_T}, h_{e_T}) \cap (D_0^-, D_0^+)$, if it is non-empty. If $b_{e_t} \leq D_0^-$ then $Ad(\mathbf{d}_{\mathcal{F}}) = (d^-, d^+)$ for $d^- = D_0^-$; otherwise $Ad(\mathbf{d}_{\mathcal{F}}) = [d^-, d^+)$ for $d^- = b_{e_T}$. By [3, Proposition 5.5], for d in this interval the derivative $\frac{\partial}{\partial d} D_R(T, \mathbf{d})$ is:

$$\cosh R \left[\sqrt{\frac{1}{\cosh^2(d/2)} - \frac{1}{\cosh^2 J(P_{v_T}(\mathbf{d}))}} - \sqrt{\frac{1}{\cosh^2(d/2)} - \frac{1}{\cosh^2 J(P_v(\mathbf{d}))}} \right]$$

Since $J(P_{v_T}(\mathbf{d})) > J(P_v(\mathbf{d}))$, this quantity is positive, and it follows that the defect sum increases with d . Therefore its minimum is at d^- .

If $d^- = b_e$ then $P_v(\mathbf{d}) \in \widetilde{\mathcal{BC}}_n$ for $\mathbf{d} = (d^-, \mathbf{d}_{\mathcal{F}})$, by Lemma 5.5, and condition (2) above holds. If $d^- = D_0^-$ then since $D_0^- \geq b_{e_T} > 0$ in this case, $P_{v_T}(\mathbf{d}) \in \widetilde{\mathcal{BC}}_{n_T}$ for $\mathbf{d} = (d^-, \mathbf{d}_{\mathcal{F}})$ as we observed above, and condition (1) holds. We also noted above that $d^- = D_0^-$ is not the longest edge of $P_{v_T}(\mathbf{d})$ in this case. \square

A two-edged tree is homeomorphic to an interval, but up to symmetry there are two possibilities for a root vertex: the intersection of the two edges, or one of the two boundary vertices. Although these two possibilities have different admissible spaces, the locations at which the associated defect function may be minimized satisfy the same criteria.

Proposition 5.10. *Let V be a graph and T a subtree of V with two edges and root vertex v_T . Let \mathcal{F} be the frontier of T in V and fix $\mathbf{d}_{\mathcal{F}} \in \mathbb{R}^{\mathcal{F}}$ with $Ad(\mathbf{d}_{\mathcal{F}}) \neq \emptyset$. For $R \geq 0$, $D_R(T, \mathbf{d})$ attains a minimum at $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ satisfying one of:*

- (1) $P_v(\mathbf{d}) \in \widetilde{\mathcal{BC}}_{n_v}$ for each $v \in T^{(0)} - \{v_T\}$, where v has valence n_v in V ; or
- (2) $P_{v_T}(\mathbf{d}) \in \widetilde{\mathcal{BC}}_{n_T}$, where v_T has valence n_T in V ; or
- (3) $J(P_v(\mathbf{d})) = J(P_{v_T}(\mathbf{d}))$ for each $v \in T^{(0)} - \{v_T\}$.

Proposition 5.10 follows directly from the two lemmas below, which separately address the possible locations for the root vertex of T .

Lemma 5.11. *Let V be a graph and T a subtree of V with two edges that share its root vertex v_T . With \mathcal{F} , $\mathbf{d}_{\mathcal{F}} \in \mathbb{R}^{\mathcal{F}}$, and $R \geq 0$ as in Proposition 5.10, its conclusions hold.*

Proof. Lemma 5.7 describes $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$ and asserts that $\sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}))$ attains a minimum somewhere on $\overline{Ad}(\mathbf{d}_{\mathcal{F}}) \times \{\mathbf{d}_{\mathcal{F}}\}$. We will show that if $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ satisfies none of the criteria of Proposition 5.10, then $\mathbf{d}_{\mathcal{E}}$ may be deformed in $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$ to reduce $\sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}))$.

Let $e_1 = e_{v_1}$ and $e_2 = e_{v_2}$, and let d_1 and d_2 be the respective lengths of their geometric duals. Then $\mathbf{d}_{\mathcal{E}} = (d_1, d_2)$. We note that as long as $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}}) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$, then reducing either of d_1 or d_2 does not increase the defect sum, since (say) $\frac{\partial}{\partial d_1} \sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}))$ is:

$$\cosh R \left[\sqrt{\frac{1}{\cosh^2(d_1/2)} - \frac{1}{\cosh^2 J(P_{v_T}(\mathbf{d}))}} - \sqrt{\frac{1}{\cosh^2(d_1/2)} - \frac{1}{\cosh^2 J(P_{v_1}(\mathbf{d}))}} \right] \geq 0$$

This follows from [3, Proposition 5.5] because $\mathbf{d}_{\mathcal{E}} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ implies that $P_{v_T}(\mathbf{d}) \in \overline{\mathcal{C}}_{n_T}$, d_1 is the largest side length of $P_{v_1}(\mathbf{d}) \in \mathcal{AC}_{n_1} - \mathcal{C}_{n_1}$, and $J(P_{v_T}(\mathbf{d})) \geq J(P_{v_1}(\mathbf{d}))$.

Now assume that $\mathbf{d}_{\mathcal{E}} = (d_1, d_2)$ does not satisfy any of criteria (1) – (3) from Proposition 5.10. Thus $P_{v_T}(\mathbf{d}) \in \mathcal{C}_n$ by (2), and by (3) we may assume that (say) $J(P_{v_T}(\mathbf{d})) > J(P_{v_1}(\mathbf{d}))$. If $P_{v_2}(\mathbf{d}) \in \overline{\mathcal{C}}_{n_2}$, then $P_{v_1}(\mathbf{d}) \notin \overline{\mathcal{C}}_{n_1}$ by (1). In this case, addressed in the paragraph below, we also have $J(P_{v_T}(\mathbf{d})) > d_2/2 = J(P_{v_2}(\mathbf{d}))$, by [3, Lemma 3.9].

Since the radius varies continuously with \mathbf{d} (see [3, Proposition 4.1]), and $P_{v_1}(\mathbf{d})$ is in the open set $\mathcal{AC}_{n-1} - \overline{\mathcal{C}}_{n-1}$, and $P_{v_T}(\mathbf{d})$ is in the open set \mathcal{C}_n there exists $\epsilon > 0$ such that for $d_1 - \epsilon < d'_1 < d_1$ and $\mathbf{d}'_{\mathcal{E}} = (d'_1, d_2)$, $P_{v_1}(\mathbf{d}') \in \mathcal{AC}_{n_1} - \overline{\mathcal{C}}_{n_1}$, $P_{v_T}(\mathbf{d}') \in \mathcal{C}_n$, and $J(P_{v_T}(\mathbf{d}')) > J(P_{v_i}(\mathbf{d}'))$ for $i = 1$ or 2 , where $\mathbf{d}' = (\mathbf{d}'_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$. Note that $P_{v_2}(\mathbf{d}') = P_{v_2}(\mathbf{d}) \in \mathcal{AC}_{n_2} - \mathcal{C}_{n_2}$. Therefore each $\mathbf{d}'_{\mathcal{E}} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$, and the defect computation above gives:

$$\sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}')) < \sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}))$$

(In particular, since $J(P_{v_T}(\mathbf{d})) > J(P_{v_1}(\mathbf{d})) > d_1/2$, the inequality is strict.)

Continuing to assume that (1) – (3) do not hold, and in particular that $J(P_{v_T}(\mathbf{d})) > J(P_{v_1}(\mathbf{d}))$, let us now suppose that $P_{v_2}(\mathbf{d}) \notin \overline{\mathcal{C}}_{n_2}$. In this case it is possible that $J(P_{v_T}(\mathbf{d})) = J(P_{v_2}(\mathbf{d}))$. We will reduce d_2 instead of d_1 , yielding $\mathbf{d}'_{\mathcal{E}} = (d_1, d'_2)$ for $d'_2 < d_2$, and $\mathbf{d}' = (\mathbf{d}'_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$. Note that [3, Lemma 4.5] implies that $\frac{\partial}{\partial d_2} J(P_{v_2}(\mathbf{d})) > \frac{1}{2} > \frac{\partial}{\partial d_2} J(P_{v_T}(\mathbf{d}))$, and indeed this estimate holds at \mathbf{d}' for as long as $P_{v_T}(\mathbf{d}') \in \mathcal{C}_{n_T}$ and $P_{v_2}(\mathbf{d}') \in \mathcal{AC}_{n_2} - \overline{\mathcal{C}}_{n_2}$.

Let $\epsilon > 0$ be small enough that if $d_2 - \epsilon < d'_2 < d_2$ and $\mathbf{d}'_{\mathcal{E}} = (d_1, d'_2)$, then $P_{v_T}(\mathbf{d}') \in \mathcal{C}_{n_T}$, $P_{v_2}(\mathbf{d}') \in \mathcal{AC}_{n_2} - \overline{\mathcal{C}}_{n_2}$ and $J(P_{v_T}(\mathbf{d}')) > J(P_{v_1}(\mathbf{d}'))$, where $\mathbf{d}' = (\mathbf{d}'_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$. By the paragraph above, $J(P_{v_T}(\mathbf{d}')) > J(P_{v_2}(\mathbf{d}'))$ for such \mathbf{d}' , and since $P_{v_1}(\mathbf{d}') = P_{v_1}(\mathbf{d}) \in \mathcal{AC}_{n_1} - \mathcal{C}_{n_1}$ it follows that $\mathbf{d}' \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$. Furthermore, the change-of-defect computation using [3, Proposition 5.5] again implies a strict decrease in defect. \square

Lemma 5.12. *Let V be a graph and T a rooted subtree with two edges, only one containing the root vertex v_T , and other vertices v_1 and v_2 . With \mathcal{F} , $\mathbf{d}_{\mathcal{F}} \in \mathbb{R}^{\mathcal{F}}$, and $R \leq \min\{d_e/2 \mid e \in \mathcal{F}\}$ as in Proposition 5.10, the conclusions of the proposition hold.*

Proof. Lemma 5.7 describes $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$ and asserts that $\sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}))$ attains a minimum somewhere on $\overline{Ad}(\mathbf{d}_{\mathcal{F}}) \times \{\mathbf{d}_{\mathcal{F}}\}$. We will show that if $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ satisfies none of the criteria above, then $\mathbf{d}_{\mathcal{E}}$ may be deformed in $\overline{Ad}(\mathbf{d}_{\mathcal{F}})$ to reduce $\sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}))$.

Take v_2 to be the opposite endpoint of the edge $e_2 = e_{v_2}$ containing v_T , let v_1 be the far endpoint of the other edge $e_1 = e_{v_1}$, and let d_1 and d_2 be the lengths of the geometric

duals to e_1 and e_2 , respectively, so that $\mathbf{d}_\mathcal{E} = (d_1, d_2)$. Assume now that $\mathbf{d} = (\mathbf{d}_\mathcal{E}, \mathbf{d}_\mathcal{F}) \in \overline{Ad}(\mathbf{d}_\mathcal{F}) \times \{\mathbf{d}_\mathcal{F}\}$ does not satisfy any of (1) – (3) from Proposition 5.10.

Since $\mathbf{d}_\mathcal{E} \in \overline{Ad}(\mathbf{d}_\mathcal{F})$ we have $J(P_{v_1}(\mathbf{d})) \leq J(P_{v_2}(\mathbf{d})) \leq J(P_{v_T}(\mathbf{d}))$. Since \mathbf{d} does not satisfy (3), at least one of these inequalities is strict. Let us suppose first that $J(P_{v_2}(\mathbf{d})) < J(P_{v_T}(\mathbf{d}))$. If $P_{v_1}(\mathbf{d}) \in \overline{\mathcal{C}}_{n_1}$, then “not (1)” implies that $P_{v_2}(\mathbf{d}) \notin \overline{\mathcal{C}}_{n_2}$, and furthermore:

$$J(P_{v_2}(\mathbf{d})) > d_2/2 > d_1/2 = J(P_{v_1}(\mathbf{d}))$$

Therefore there exists $\epsilon > 0$ such that for all d'_2 with $d_2 - \epsilon < d'_2 < d_2$, taking $\mathbf{d}'_\mathcal{E} = (d_1, d'_2)$ and $\mathbf{d}' = (\mathbf{d}'_\mathcal{E}, \mathbf{d}_\mathcal{F})$ we have $P_{v_2}(\mathbf{d}') \notin \overline{\mathcal{C}}_{n_2}$, $P_{v_T}(\mathbf{d}') \in \mathcal{C}_n$, and $J(P_{v_T}(\mathbf{d}')) > J(P_{v_2}(\mathbf{d}'))$. We note that $P_{v_1}(\mathbf{d}') = P_{v_1}(\mathbf{d})$ for all such \mathbf{d}' . [3, Proposition 5.5] implies that $\frac{\partial}{\partial d_2} \sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}))$ is:

$$\cosh R \left[\sqrt{\frac{1}{\cosh^2(d_2/2)} - \frac{1}{\cosh^2 J(P_{v_T}(\mathbf{d}))}} - \sqrt{\frac{1}{\cosh^2(d_2/2)} - \frac{1}{\cosh^2 J(P_{v_2}(\mathbf{d}))}} \right]$$

As long as $J(P_{v_T}(\mathbf{d})) > J(P_{v_2}(\mathbf{d}))$, this quantity is positive, so decreasing d_2 decreases the defect sum. Thus with \mathbf{d}' as above we have $\sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}')) < \sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}))$.

Continuing to suppose that $J(P_{v_2}(\mathbf{d})) < J(P_{v_T}(\mathbf{d}))$, let us now also assume that $P_{v_1}(\mathbf{d}) \notin \overline{\mathcal{C}}_{n_1}$. [3, Lemma 4.5] implies that decreasing d_1 has the effect of decreasing $J(P_{v_1}(\mathbf{d}))$ but increasing $J(P_{v_2}(\mathbf{d}))$, since $P_{v_2}(\mathbf{d}) \in \mathcal{AC}_{n_2} - \mathcal{C}_{n_2}$ has longest side d_2 . Thus there exists $\epsilon > 0$ such that for $d_1 - \epsilon < d'_1 < d_1$, taking $\mathbf{d}'_\mathcal{E} = (d'_1, d_2)$ and $\mathbf{d}' = (\mathbf{d}'_\mathcal{E}, \mathbf{d}_\mathcal{F})$ we have $P_{v_1}(\mathbf{d}') \notin \overline{\mathcal{C}}_{n_1}$, $P_{v_2}(\mathbf{d}') \in \mathcal{AC}_{n_2} - \mathcal{C}_{n_2}$, and $J(P_{v_2}(\mathbf{d}')) < J(P_{v_T}(\mathbf{d}'))$. Furthermore, $P_{v_T}(\mathbf{d}') = P_{v_T}(\mathbf{d}) \in \mathcal{C}_n$ for all such \mathbf{d}' , so $\frac{\partial}{\partial d_1} \sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}'))$ is:

$$\cosh R \left[\sqrt{\frac{1}{\cosh^2(d'_1/2)} - \frac{1}{\cosh^2 J(P_{v_2}(\mathbf{d}'))}} - \sqrt{\frac{1}{\cosh^2(d'_1/2)} - \frac{1}{\cosh^2 J(P_{v_1}(\mathbf{d}'))}} > 0 \right]$$

Thus we again find that $\sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}')) < \sum_{v \in T^{(0)}} D_R(P_v(\mathbf{d}))$ for $d'_1 < d_1$. (Note that even if $J(P_{v_1}(\mathbf{d})) = J(P_{v_2}(\mathbf{d}))$, strict inequality holds for \mathbf{d}' by the above, and so the strict inequality of defect sums is also accurate.)

Let us finally suppose that $J(P_{v_2}(\mathbf{d})) = J(P_{v_T}(\mathbf{d}))$. Then since (3) does not hold, $J(P_{v_1}(\mathbf{d})) < J(P_{v_2}(\mathbf{d}))$. Since (2) does not hold we have $P_{v_T}(\mathbf{d}) \in \mathcal{C}_n$, so $J(P_{v_T}(\mathbf{d})) = J(P_{v_2}(\mathbf{d})) > d_2/2$ and so also $P_{v_2}(\mathbf{d}) \notin \overline{\mathcal{C}}_{n_2}$. [3, Lemma 4.5] implies that reducing d_2 reduces the radius of $P_{v_2}(\mathbf{d})$ faster than that of $P_{v_T}(\mathbf{d})$, and it follows that d_2 may be reduced slightly keeping $\mathbf{d}_\mathcal{E} \in \overline{Ad}(\mathbf{d}_\mathcal{F})$. A derivative computation as above shows that this reduces the defect. \square

6. DEFECT BOUNDS FROM SIDE LENGTH BOUNDS

Definition 6.1. For \mathcal{F} finite and $\mathbf{b}_\mathcal{F}, \mathbf{d}_\mathcal{F} \in \mathbb{R}^\mathcal{F}$, say $\mathbf{d}_\mathcal{F} \geq \mathbf{b}_\mathcal{F}$ if $d_f \geq b_f$ for each $f \in \mathcal{F}$, where $\mathbf{b}_\mathcal{F} = (b_f \mid f \in \mathcal{F})$ and $\mathbf{d}_\mathcal{F} = (d_f \mid f \in \mathcal{F})$.

This section describes an algorithm with the following:

Input: A rooted tree T with frontier \mathcal{F} , $R \geq 0$, and $\mathbf{b}_\mathcal{F} \in (\mathbb{R}^+)^{\mathcal{F}}$.

Output: $D > 0$ such that $D_R(T, (\mathbf{d}_\mathcal{E}, \mathbf{d}_\mathcal{F})) \geq D$ for all $\mathbf{d}_\mathcal{F} \geq \mathbf{b}_\mathcal{F}$ and $\mathbf{d}_\mathcal{E} \in \overline{Ad}(\mathbf{d}_\mathcal{F})$.

We begin with some *a priori* bounds.

Lemma 6.2. *Let T be a rooted tree with root vertex v_T , edge set \mathcal{E} , and frontier \mathcal{F} such that each $v \in T^{(0)}$ is at least three-valent in $T \cup \bigcup \{e \mid (e, v) \in \mathcal{F}\}$. Fix $\mathbf{b}_{\mathcal{F}} \in \mathbb{R}^{\mathcal{F}}$. For $v \in T^{(0)} - \{v_T\}$ let e_0, \dots, e_{n-1} be the set of edges containing v , with $e_v = e_0$, and define $P_v^h(\mathbf{b}_{\mathcal{F}}) = (h_0(b_1, \dots, b_{n-1}), b_1, \dots, b_{n-1}) \in \widetilde{\mathcal{HC}}_n$, where h_0 is as in [3, Lemma 3.10] and:*

$$b_i = \begin{cases} b_{e_i}(\mathbf{b}_{\mathcal{F}}) & e_i \in \mathcal{E} \\ b_{e_i} & e_i \in \mathcal{F} \end{cases}$$

Then for $R \geq 0$, $\mathbf{d}_{\mathcal{F}} \geq \mathbf{b}_{\mathcal{F}}$, $\mathbf{d}_{\mathcal{E}} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ and $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$, $D_R(P_v(\mathbf{d})) > D_R(P_v^h(\mathbf{b}_{\mathcal{F}}))$, where $P_v(\mathbf{d})$ is as in Definition 5.2.

Proof. Fix $\mathbf{d}_{\mathcal{F}} \geq \mathbf{b}_{\mathcal{F}}$ and $\mathbf{d}_{\mathcal{E}} = (d_e \mid e \in \mathcal{E}) \in Ad(\mathbf{d}_{\mathcal{F}})$. By Lemma 5.5, $d_e \geq b_e(\mathbf{d}_{\mathcal{F}}) \geq b_e(\mathbf{b}_{\mathcal{F}})$ for each $e \in \mathcal{E}$. Thus for $v \in T^{(0)} - \{v_T\}$ and e_0, \dots, e_{n-1} as described in the hypotheses above, $b_i \leq d_{e_i}$ for each $i > 0$. Criterion (1) of Lemma 5.7 implies that $P_v(\mathbf{d}) \in \mathcal{AC}_n - \mathcal{C}_n$ has longest edge d_{e_v} , so [3, Corollary 5.11] implies $D_R(P_v(\mathbf{d})) > D_R(P_v^h(\mathbf{b}_{\mathcal{F}}))$. \square

Proposition 6.3. *Let T be a rooted tree with root vertex v_T , edge set \mathcal{E} , and frontier \mathcal{F} such that each $v \in T^{(0)}$ is at least three-valent in $T \cup \bigcup \{e \mid (e, v) \in \mathcal{F}\}$. Fix $\mathbf{b}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$. For a subtree T_0 of T with $v_T \in T_0$, let \mathcal{E}_0 and \mathcal{F}_0 be the edge set and frontier (in $T \cup \bigcup \{e \mid (e, v) \in \mathcal{F}\}$) of T_0 , respectively, and define $\mathbf{b}_{\mathcal{F}_0} = (b_e \mid e \in \mathcal{F}_0)$ by:*

$$b_e = \begin{cases} b_e & e \in \mathcal{F} \\ b_e(\mathbf{b}_{\mathcal{F}}) & e \in \mathcal{E} \end{cases}$$

Let $D_0 = \inf\{D_R(T_0, (\mathbf{d}_{\mathcal{E}_0}, \mathbf{d}_{\mathcal{F}_0})) \mid \mathbf{d}_{\mathcal{E}_0} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}_0}), \mathbf{d}_{\mathcal{F}_0} \geq \mathbf{b}_{\mathcal{F}_0}\}$. Then for any $\mathbf{d}_{\mathcal{F}} \geq \mathbf{b}_{\mathcal{F}}$ and $\mathbf{d}_{\mathcal{E}} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$,

$$D_R(T, \mathbf{d}_{\mathcal{F}}) > D_0 + \sum_{v \in T^{(0)} - T_0^{(0)}} D_R(P_v^h(\mathbf{b}_{\mathcal{F}})),$$

where $P_v^h(\mathbf{b}_{\mathcal{F}})$ is as in Lemma 6.2.

Proof. A fixed pair $\mathbf{d}_{\mathcal{F}} = (d_e \mid e \in \mathcal{F}) \geq \mathbf{b}_{\mathcal{F}}$ and $\mathbf{d}_{\mathcal{E}} = (d_e \mid e \in \mathcal{E}) \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ determines $\mathbf{d}_{\mathcal{F}_0}$ and $\mathbf{d}_{\mathcal{E}_0}$ simply by taking the appropriate entries of $\mathbf{d}_{\mathcal{F}}$ or $\mathbf{d}_{\mathcal{E}}$. Lemma 5.5 and the construction of $\mathbf{b}_{\mathcal{F}_0}$ then imply that $\mathbf{d}_{\mathcal{F}_0} \geq \mathbf{b}_{\mathcal{F}_0}$.

Taking $\mathbf{d}_0 = (\mathbf{d}_{\mathcal{E}_0}, \mathbf{d}_{\mathcal{F}_0})$, it is clear by Definition 5.2 and the construction of $\mathbf{d}_{\mathcal{E}_0}$ and $\mathbf{d}_{\mathcal{F}_0}$ that for each $v \in T_0$, $P_v(T_0, \mathbf{d}_0) = P_v(T, \mathbf{d})$, where $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$. Definition 5.2 thus implies:

$$D_R(T, \mathbf{d}) = D_R(T_0, \mathbf{d}_0) + \sum_{v \in T^{(0)} - T_0^{(0)}} D_R(P_v(\mathbf{d}))$$

Lemma 5.7 implies that $\mathbf{d}_{\mathcal{E}_0} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}_0})$, since $\mathbf{d}_{\mathcal{E}} \in \mathbf{d}_{\mathcal{F}}$ by hypothesis, so $D_R(T_0, \mathbf{d}_0) \geq d_0$ and the result follows from Lemma 6.2. \square

Proposition 6.3 can be used in conjunction with the result below to give *a priori* bounds.

Lemma 6.4. *Let T be a rooted tree with root vertex v_T , edge set \mathcal{E} , and frontier \mathcal{F} such that each $v \in T^{(0)}$ is at least three-valent in $T \cup \bigcup \{e \mid (e, v) \in \mathcal{F}\}$. For $\mathbf{b}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$, let $\mathbf{b}_{\mathcal{E}} = (b_e(\mathbf{b}_{\mathcal{F}}) \mid e \in \mathcal{E}) \in \mathbb{R}^{\mathcal{E}}$ and take $\mathbf{b} = (\mathbf{b}_{\mathcal{E}}, \mathbf{b}_{\mathcal{F}})$. Enumerate the edges of V containing v_T as e_0, \dots, e_{n-1} so that b_{e_0} is largest. With $b_0: (\mathbb{R}^+)^{n-1} \rightarrow \mathbb{R}^+$ as in [3, Lemma 3.4], define*

$$B_{e_0} = \begin{cases} b_0(b_{e_1}, \dots, b_{e_{n-1}}) & \text{if } b_{e_0} > b_0(b_{e_1}, \dots, b_{e_{n-1}}) \\ b_{e_0} & \text{otherwise} \end{cases}$$

For $R \geq 0$, let $M_R(v_T, \mathbf{b}_{\mathcal{F}}) = D_R(B_{e_0}, b_{e_1}, \dots, b_{e_{n-1}})$. Then $D_R(P_{v_T}(\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})) \geq M_R(v_T, \mathbf{b}_{\mathcal{F}})$ for each $\mathbf{d}_{\mathcal{F}} \geq \mathbf{b}_{\mathcal{F}}$, and $\mathbf{d}_{\mathcal{E}} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$.

Proof. Given $\mathbf{d}_{\mathcal{E}}$ and $\mathbf{d}_{\mathcal{F}}$ as above, Lemma 5.5 implies that $d_{e_i} \geq b_{e_i}$ for each $i \in \{0, \dots, n-1\}$, and therefore also that $d_{e_0} \geq B_{e_0}$. Since $P_{v_T}(\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}}) \in \widetilde{\mathcal{C}}_n \cup \widetilde{\mathcal{BC}}_n$ by Lemma 5.7, and $(B_{e_0}, b_{e_1}, \dots, b_{e_{n-1}}) \in \widetilde{\mathcal{C}}_n \cup \widetilde{\mathcal{BC}}_n$ by construction, [3, Lemma 6.6] implies the result. \square

Remark 6.5. With the hypotheses of Lemma 6.4, if v_T is three-valent in V and $b_{e_0} > b_0(b_{e_1}, b_{e_2})$, the conclusion may be improved, using [3, Lemma 6.9], taking:

$$M_R(v_T, \mathbf{b}_{\mathcal{F}}) = \min\{D_R(b_{e_0}, b_{e_1}, b'_2), D_R(b_{e_0}, b'_1, b_{e_2})\},$$

where $\cosh b_{e_0} = \cosh b'_1 + \cosh b_{e_2} - 1 = \cosh b_{e_1} + \cosh b'_2 - 1$.

Given a rooted tree T with root vertex v_T , edge set \mathcal{E} , and frontier \mathcal{F} such that each vertex of T is at least three-valent in $T \cup \bigcup_{e \in \mathcal{F}} e$, for $\mathbf{b}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ and $R \geq 0$, the procedure below can be algorithmically implemented:

- (1) For each $e \in \mathcal{E}$, compute $b_e(\mathbf{b}_{\mathcal{F}})$ as in Lemma 5.5 (see Remark 5.6).
- (2) For each $v \in T^{(0)} - \{v_T\}$, compute $D_R(P_v^h(\mathbf{b}_{\mathcal{F}}))$ for $P_v^h(\mathbf{b}_{\mathcal{F}})$ as in Lemma 6.2.
- (3) Compute $M_R(v_T, \mathbf{b}_{\mathcal{F}})$ with Lemma 6.4, or if v_T has valence three, with Remark 6.5.
- (4) Let $D = M_R(v_T, \mathbf{b}_{\mathcal{F}}) + \sum_{v \in T^{(0)} - \{v_T\}} D_R(P_v^h(\mathbf{b}_{\mathcal{F}}))$.

By Proposition 6.3 (taking $T_0 = \{v_T\}$) and Lemma 6.4, D as defined above is a lower bound on $D_R(T, (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}}))$ for any $\mathbf{d}_{\mathcal{F}} \geq \mathbf{b}_{\mathcal{F}}$, and $\mathbf{d}_{\mathcal{E}} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$

Below we will describe how to improve the procedure above under the assumption that $D_R(T, \mathbf{d})$ attains its minimum at a point of $Ad(\mathbf{d}_{\mathcal{F}})$ satisfying one of the three criteria of Proposition 5.10. We will treat these cases separately.

6.1. Case (1): $P_v(\mathbf{d}) \in \widetilde{\mathcal{BC}}_{n_v}$ for all $v \in T^{(0)} - \{v_T\}$. Lemma 5.5 implies that each $\mathbf{d}_{\mathcal{F}} \in \mathbb{R}^{\mathcal{F}}$ determines a unique $\mathbf{d}_{\mathcal{E}} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ such that $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ falls into this case. For such \mathbf{d} , [3, Corollary 5.8] implies:

Lemma 6.6. *Let T be a rooted tree with root vertex v_T , edge set \mathcal{E} , and frontier \mathcal{F} such that each $v \in T^{(0)}$ is at least three-valent in $T \cup \bigcup \{e \mid (e, v) \in \mathcal{F}\}$. Fix $\mathbf{b}_{\mathcal{F}} \in \mathbb{R}^{\mathcal{F}}$, let $\mathbf{b}_{\mathcal{E}} = (b_e(\mathbf{b}_{\mathcal{F}}) \mid e \in \mathcal{E})$, and for $v \in T^{(0)} - \{v_T\}$ define $P_v^b(\mathbf{b}_{\mathcal{F}}) = P_v(\mathbf{b}_{\mathcal{E}}, \mathbf{b}_{\mathcal{F}}) \in \widetilde{\mathcal{BC}}_{n_v}$. Then for $R \geq 0$, $\mathbf{d}_{\mathcal{F}} \geq \mathbf{b}_{\mathcal{F}}$, $\mathbf{d}_{\mathcal{E}} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ and $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$ such that $P_v(\mathbf{d}) \in \widetilde{\mathcal{BC}}_{n_v}$ for all $v \in T^{(0)} - \{v_T\}$, $D_R(P_v(\mathbf{d})) > D_R(P_v^b(\mathbf{b}_{\mathcal{F}}))$ for each such v .*

Using Lemma 6.6, we may improve the basic algorithm in this case by replacing the computation of $D_R(P_v^h(\mathbf{b}_\mathcal{F}))$ in step (2) with that of $D_R(P_v^b(\mathbf{b}_\mathcal{F}))$, and in step (4) taking $D = M_R(v_T, \mathbf{b}_\mathcal{F}) + \sum_{v \in T^{(0)} - \{v_T\}} D_R(P_v^b(\mathbf{b}_\mathcal{F}))$.

6.2. Case (2): $P_{v_T}(\mathbf{d}) \in \widetilde{\mathcal{BC}}_{n_T}$. Our main advantage in this case is the following improved version of Lemma 6.4.

Lemma 6.7. *Let T be a rooted tree with root vertex v_T , edge set \mathcal{E} , and frontier \mathcal{F} such that each $v \in T^{(0)}$ is at least three-valent in $T \cup \bigcup \{e \mid (e, v) \in \mathcal{F}\}$. For $\mathbf{b}_\mathcal{F} \in (\mathbb{R}^+)^{\mathcal{F}}$, let $\mathbf{b}_\mathcal{E} = (b_e(\mathbf{b}_\mathcal{F}) \mid e \in \mathcal{E}) \in \mathbb{R}^\mathcal{E}$ and take $\mathbf{b} = (\mathbf{b}_\mathcal{E}, \mathbf{b}_\mathcal{F})$, and enumerate the edges of V containing v_T as e_0, \dots, e_{n-1} . With $b_0: (\mathbb{R}^+)^{n-1} \rightarrow \mathbb{R}^+$ as in [3, Lemma 3.4], for each i define*

$$B_{e_i} = b_0(b_{e_0}, \dots, \hat{b}_{e_i}, \dots, b_{e_{n-1}})$$

Then for $R \geq 0$, $\mathbf{d}_\mathcal{F} \geq \mathbf{b}_\mathcal{F}$, and $\mathbf{d}_\mathcal{E} \in \overline{Ad}(\mathbf{d}_\mathcal{F})$ such that $P_{v_T}(\mathbf{d}_\mathcal{E}, \mathbf{d}_\mathcal{F}) \in \widetilde{\mathcal{BC}}_{n_T}$ has longest side dual to e_i , $D_R(P_{v_T}(\mathbf{d}_\mathcal{E}, \mathbf{d}_\mathcal{F})) \geq D_R(b_{e_0}, \dots, B_{e_i}, \dots, b_{e_{n-1}})$.

This follows from Lemma 5.5 and [3, Corollary 5.8] as in the proof of Lemma 6.4. If the longest edge of $P_{v_T}(\mathbf{d})$ is dual to an element of \mathcal{E} , we may further augment Lemma 6.4:

Lemma 6.8. *Let T be a rooted tree with root vertex v_T , edge set \mathcal{E} , and frontier \mathcal{F} such that each $v \in T^{(0)}$ is at least three-valent in $T \cup \bigcup \{e \mid (e, v) \in \mathcal{F}\}$. For $\mathbf{d}_\mathcal{F} \in (\mathbb{R}^+)^{\mathcal{F}}$ and $\mathbf{d}_\mathcal{E} \in \overline{Ad}(\mathbf{d}_\mathcal{F})$ such that $P_{v_T}(\mathbf{d}_\mathcal{E}, \mathbf{d}_\mathcal{F}) \in \widetilde{\mathcal{BC}}_{n_T}$ has longest side dual to $e \in \mathcal{E}$, $P_v(\mathbf{d}_\mathcal{E}, \mathbf{d}_\mathcal{F}) \in \widetilde{\mathcal{BC}}_{n_v}$, where v is the initial vertex of e .*

Proof. Let $\mathbf{d} = (\mathbf{d}_\mathcal{E}, \mathbf{d}_\mathcal{F})$. Since $P_{v_T}(\mathbf{d}) \in \widetilde{\mathcal{BC}}_{n_T}$ by hypothesis, [3, Lemma 3.9] implies that $J(P_{v_T}(\mathbf{d})) = d_i/2$, where d_i is the length of the geometric dual to e_i . On the other hand, $P_{v_i}(\mathbf{d})$ also has longest side dual to e_i by Lemma 5.7, which further implies that $J(P_{v_i}(\mathbf{d})) \leq J(P_{v_T}(\mathbf{d}))$. Since $J(P_{v_0}(\mathbf{d})) \geq d_i/2$ by [3, Lemma 3.6], it is equal to $d_i/2$, and therefore the Lemma holds by [3, Lemma 3.9]. \square

From Lemma 5.5 and [3, Corollary 5.8] we thus directly obtain:

Corollary 6.9. *With the hypotheses of Lemma 6.7, for $R \geq 0$, $\mathbf{d}_\mathcal{F} \geq \mathbf{b}_\mathcal{F}$, and $\mathbf{d}_\mathcal{E} \in \overline{Ad}(\mathbf{d}_\mathcal{F})$ such that $P_{v_T}(\mathbf{d}_\mathcal{E}, \mathbf{d}_\mathcal{F}) \in \widetilde{\mathcal{BC}}_{n_T}$ has longest side dual to $e_i \in \mathcal{E}$, $D_R(P_{v_i}(\mathbf{d})) \geq D_R(P_{v_i}^b(\mathbf{b}_\mathcal{F}))$, where $\mathbf{d} = (\mathbf{d}_\mathcal{E}, \mathbf{d}_\mathcal{F})$, v_i is the initial vertex of e_i , and $P_{v_i}^b(\mathbf{b}_\mathcal{F})$ is as in Lemma 6.6.*

Remark 6.10. With the hypotheses of Corollary 6.9, if v_i is trivalent in $T \cup \bigcup \{e \mid (e, v) \in \mathcal{F}\}$ and $b_{e_i}(\mathbf{b}_\mathcal{F}) < B_{e_i}$, then using [3, Lemma 6.9] as in Remark 6.5 we have:

$$D_R(P_{v_i}(\mathbf{d})) \geq \min\{D_R(B_{e_i}, b_{f_1}, b'_{f_2}), D_R(b_{e_0}, b'_{f_1}, b_{f_2})\}$$

Here f_1 and f_2 are the other edges containing v_i , b_{f_1} and b_{f_2} are as in Lemma 6.6, and $\cosh B_{e_i} = \cosh b'_{f_1} + \cosh b_{f_2} - 1 = \cosh b_{f_1} + \cosh b'_{f_2} - 1$.

In order to improve the basic algorithm in this case, enumerate the edges containing v_T as e_0, \dots, e_{n-1} , and in step (3) of the basic algorithm replace the computation of $M_R(v_T, \mathbf{b}_\mathcal{F})$ with those of $M_R^{(i)}(v_T, \mathbf{b}_\mathcal{F}) = D_R(b_{e_0}, \dots, B_{e_i}, \dots, b_{e_{n-1}})$ for each i , where B_{e_i} is as in Lemma 6.7. It is useful now to divide into two subcases:

Case (2)(A): In step (4) of the basic algorithm, replace D with:

$$D_A = \min \left\{ M_R^{(i)}(v_T, \mathbf{b}_{\mathcal{F}}) \mid e_i \in \mathcal{F} \right\} + \sum_{v \in T^{(0)} - \{v_T\}} D_R(P_v^h(\mathbf{b}_{\mathcal{F}}))$$

Case (2)(B): In step (2) of the basic algorithm, also compute $D_R(P_{v_i}^b(\mathbf{b}_{\mathcal{F}}))$ for each i such that $e_i \in \mathcal{E}$, where v_i is the initial vertex of e_i , and in step (4) replace D with:

$$D_B = \min_{e_i \in \mathcal{E}} \left\{ M_R^{(i)}(v_T, \mathbf{b}_{\mathcal{F}}) + D_R(P_{v_i}^b(\mathbf{b}_{\mathcal{F}})) + \sum_{v \in T^{(0)} - \{v_T, v_i\}} D_R(P_v^h(\mathbf{b}_{\mathcal{F}})) \right\}$$

For each i such that v_i is trivalent, $D_R(P_{v_i}^b(\mathbf{b}_{\mathcal{F}}))$ can be replaced by the computation from Remark 6.10 if $b_{e_i}(\mathbf{b}_{\mathcal{F}}) < B_{e_i}$.

By the results above, $D = \min\{D_A, D_B\}$ bounds $D_R(T, (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}}))$ below for any $\mathbf{d}_{\mathcal{F}} \geq \mathbf{b}_{\mathcal{F}}$ and $\mathbf{d}_{\mathcal{E}} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$.

6.3. Case (3): $J(P_v(\mathbf{d})) = J(P_{v_T}(\mathbf{d}))$ for all $v \in T^{(0)} - \{v_T\}$. Here we have:

Lemma 6.11. *Let T be a rooted tree with root vertex v_T , edge set \mathcal{E} , and frontier \mathcal{F} such that each $v \in T^{(0)}$ is at least three-valent in $T \cup \bigcup \{e \mid (e, v) \in \mathcal{F}\}$. For $\mathbf{d}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$ and $\mathbf{d}_{\mathcal{E}} \in \overline{Ad}(\mathbf{d}_{\mathcal{F}})$ such that $J(P_v(\mathbf{d})) = J(P_{v_T}(\mathbf{d}))$ for each $v \in T^{(0)} - \{v_T\}$, where $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$, there exists $P \in \tilde{\mathcal{C}}_{|\mathcal{F}|} \cup \widetilde{\mathcal{BC}}_{|\mathcal{F}|}$ such that $D_R(T, \mathbf{d}) = D_R(P)$ for any $R \geq 0$.*

Proof. For $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$, note that $P_{v_T}(\mathbf{d}) \in \tilde{\mathcal{C}}_{n_T} \cup \widetilde{\mathcal{BC}}_{n_T}$ by Lemma 5.7. Now fix a subtree T_0 of T with $v_T \in T_0$ and $T = T_0 \cup e_0$ for some $e_0 \in \mathcal{E}$, and assume that the following holds for T_0 : there is a cyclic polygon P_0 in \mathbb{H}^2 that contains a copy of $P_v(\mathbf{d})$ for each $v \in T_0^{(0)}$, such that $P_0 = \bigcup_{v \in T_0^{(0)}} P_v(\mathbf{d})$ and $P_v(\mathbf{d}) \cap P_w(\mathbf{d})$ contains more than one point if and only if v and w bound an edge e of T_0 , in which case $P_v(\mathbf{d}) \cap P_w(\mathbf{d})$ is the geometric dual to e .

The edge set of P_0 is in one-to-one correspondence with the frontier \mathcal{F}_0 of T_0 in V , and we will assume that P_0 has the same center x_0 and radius J as $P_v(\mathbf{d}) \subset P_0$ for each $v \in T_0^{(0)}$. Thus in particular, $P_0 \in \tilde{\mathcal{C}}_{|\mathcal{F}_0|} \cup \widetilde{\mathcal{BC}}_{|\mathcal{F}_0|}$ by [3, Lemma 3.9], since this implies that $x_0 \in P_{v_T}(\mathbf{d}) \subset P_0$.

Let $\{v_0\} = T^{(0)} - T_0^{(0)}$, and enumerate the edges of V containing v_0 as e_0, \dots, e_{n-1} , where e_0 is described above. Then $e_i \in \mathcal{F}$ for each $i > 0$, since $T = T_0 \cup e_0$, and $\mathcal{F}_0 = \{e_0\} \cup (\mathcal{F} - \{e_i\}_{i=1}^{n-1})$. Since e_0 is necessarily the first edge of the path in T joining v_0 to v_T , v_0 is its initial vertex with the orientation from Lemma 3.7.

Since $e_0 \in \mathcal{F}_0$, P_0 has an edge corresponding to its geometric dual γ_0 . Arrange a copy of $P_{v_0}(\mathbf{d})$ so that it intersects P_0 in γ_0 . The isosceles triangle T_0 determined by γ_0 and the center of $P_{v_0}(\mathbf{d})$ has equal sides of length $J(P_{v_0}(\mathbf{d})) = J(P_{v_T}(\mathbf{d})) = J$ by hypothesis. Furthermore, Lemma 5.7 implies that $P_{v_0}(\mathbf{d}) \in \widetilde{\mathcal{AC}}_n - \tilde{\mathcal{C}}_n$ has longest side γ_0 , so $\gamma_0 = T_0 \cap P_{v_0}(\mathbf{d})$ by [3, Lemma 1.6]. It follows that T_0 intersects the interior of P_0 .

On the other hand, the triangle determined by γ_0 and the center x_0 of P_0 has two sides of length J and by [3, Lemma 1.6] is contained in P_0 , since $P_0 \in \tilde{\mathcal{C}}_{|\mathcal{F}_0|} \cup \widetilde{\mathcal{BC}}_{|\mathcal{F}_0|}$. Since this triangle has the same side length collection as T_0 , share γ_0 with it, and is on the same side of γ_0 it is

identical to T_0 . Therefore x_0 is the center of $P_{v_0}(\mathbf{d})$, so by [3, Lemma 1.4], $P = P_0 \cup P_{v_0}(\mathbf{d})$ is a cyclic polygon with center x_0 and radius J .

If w_0 is the terminal endpoint of e_0 , then $P_{w_0}(\mathbf{d}) \subset P_0$ contains γ_0 , so $\gamma_0 = P_{v_0}(\mathbf{d}) \cap P_{w_0}(\mathbf{d})$ and P satisfies the hypotheses for T that P_0 satisfied for T_0 . It is easy to see that $D_R(P) = D_R(P_0) + D_R(P_{v_0}(\mathbf{d}))$, so the result follows by an inductive argument. \square

The corollary below thus follows directly from [3, Corollary 5.8], and supplies the required lower bound without appeal to the basic algorithm.

Corollary 6.12. *Let T be a rooted tree with root vertex v_T , edge set \mathcal{E} , and frontier \mathcal{F} such that each $v \in T^{(0)}$ is at least three-valent in $T \cup \bigcup \{e \mid (e, v) \in \mathcal{F}\}$. Fix $\mathbf{b}_{\mathcal{F}} \in (\mathbb{R}^+)^{\mathcal{F}}$, and enumerate the edges in \mathcal{F} as $\{e_0, \dots, e_{n-1}\}$ so that b_{e_0} is maximal. Define:*

$$B_{e_0} = \begin{cases} b_0(b_{e_1}, \dots, b_{e_{n-1}}) & \text{if } b_{e_0} > b_0(b_{e_1}, \dots, b_{e_{n-1}}) \\ b_{e_0} & \text{otherwise} \end{cases}$$

Then for $R \geq 0$, $\mathbf{d}_{\mathcal{F}} \geq \mathbf{b}_{\mathcal{F}}$, $\mathbf{d}_{\mathcal{F}} \geq \mathbf{b}_{\mathcal{F}}$, and $\mathbf{d}_{\mathcal{E}} \in \overline{\text{Ad}}(\mathbf{d}_{\mathcal{F}})$ such that $J(P_v(\mathbf{d})) = J(P_{v_T}(\mathbf{d}))$ for each $v \in T^{(0)} - \{v_T\}$, where $\mathbf{d} = (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})$, $D_R(T, \mathbf{d}) \geq D_R(B_{e_0}, b_{e_1}, \dots, b_{e_{n-1}})$.

7. COMPUTATIONS

This section is devoted to applying our previous results to prove Theorem 0.1. For r_{β} as described in the theorem, $2.8298 < \cosh r_{\beta} < 2.8299$, and if $d_{\beta} = 2r_{\beta}$ then $15.0166 < \cosh d_{\beta} < 15.0167$. Let r_1 and d_1 satisfy $\cosh r_1 = 2.8298$ and $\cosh d_1 = 15.0166$, respectively. Then $r_1 < r_{\beta}$ and $d_1 < d_{\beta}$, and it is easy to show that $d_1 > 2r_1$.

Table 1 records the radius- r_1 defect of the symmetric, centered n -gons $P_n(d_1)$ for $n = 3$ through 6. These computations use [3, Lemma 6.6]. In each case we have truncated the result after five decimal places, so the actual defect value is greater than what is displayed.

n	$D_{r_1}(P_n(d_1))$	n	$D_{r_1}(P_n(d_1))$	n	$D_{r_1}(P_n(d_1))$	n	$D_{r_1}(P_n(d_1))$
3	0.12586	4	0.56593	5	1.22041	6	2.00496

TABLE 1. Radius- r_1 defects of highly symmetric polygons.

We will perform an analogous computation for centered dual 2-cells, but initially focus only on those with five frontier edges. To begin, let us note that an Euler characteristic computation implies:

Remark 7.1. If T is a tree with edge set \mathcal{E} and frontier \mathcal{F} , such that each vertex of T is at least three-valent in $T \cup \{e \mid (e, v) \in \mathcal{F} \text{ for some } v \in T^{(0)}\}$, then $|\mathcal{F}| \geq |\mathcal{E}| + 3$.

This implies in particular that for a 2-cell Q of the centered dual decomposition containing a component T of $V_n^{(1)}$, if Q has five edges then T has at most two. Carrying the same argument further, we find that the possibilities for such Q are exactly those showed in Figure 7.1. In the figure, Voronoi edges are dashed and black, and centered dual edges are solid and red.




FIGURE 7.1. Possibilities for (non-Delaunay) 5-edged centered dual 2-cells.

We have labeled the possibilities by the corresponding components of $V_n^{(1)}$ (in bold), where subscripts describe the valence of each vertex, with the root vertex in bold.

	Basic	Case (1)	Case (2)(A)	Case (2)(B)	Case (3)
$D_{r_1}(T_{3,4}, \mathbf{b})$	1.00510	1.17816	1.57569	N/A	N/A
$D_{r_1}(T_{4,3}, \mathbf{b})$	1.63705	1.77971	1.71113	N/A	N/A
$D_{r_1}(T_{3,3,3}, \mathbf{b})$	0.80915	1.15527	1.28432	1.38738	1.22041
$D_{r_1}(T_{3,3,3}, \mathbf{b})$	1.24735	1.56044	1.38585	1.38738	1.22041

 TABLE 2. $D_{r_1}(T, \mathbf{b})$ for the T from Figure 7.1, where $\mathbf{b} = (d_1, d_1, d_1, d_1, d_1)$.

Table 2 records the output of computer programs implementing the algorithms of Section 6 for the trees of Figure 7.1. (The programs are in the supplementary materials.) Each tree falls under the purview of the improvements to the basic algorithm described in the second half of the section, by Lemma 5.9 for the one-edged trees and Proposition 5.10 for the others. Since Lemma 5.9 does not allow Case (2)(B) or (3) of Proposition 5.10, we placed “N/A”s in the corresponding table entries. We have boxed the best bound for each tree: the basic algorithm’s output or the minimum of the improvements’ (whichever is larger).

Corollary 7.2. *If a closed orientable hyperbolic surface F of genus 2 has injectivity radius at least $d_1/2$, where $\cosh d_1 = 15.0166$, at $x \in F$, then no two-cell of the centered dual tessellation of F determined by x has more than four edges.*

Proof. If F has injectivity radius at least $d_1/2$ at x , then since $r_1 < d_1/2$ a hyperbolic disk U of radius r_1 is embedded in F , centered at x . The area of U is $2\pi(\cosh r_1 - 1) > 2\pi \cdot 1.8298$, so the area of its complement in F is less than $2\pi \cdot 0.1702 < 1.07$.

For each 2-cell P of the centered dual decomposition of F determined by $\{x\}$, Proposition 4.2 implies that $D_{r_1}(P)$ is the area of $P - (P \cap U)$. Let us first assume that P has five edges. Each is a geodesic arc that begins and ends at x , so its length is at least d_1 . If P is centered, then [3, Corollary 5.8] implies that $D_{r_1}(P) > D_{r_1}(P_n(d_1)) > 1.22$, by Table 1. This contradicts the fact that the total area complementary to U in F is less than 1.07.

If P contains a component T of $V_n^{(1)}$, where V is the Voronoi tessellation determined by $\{x\}$, then T is one of the possibilities pictured in Figure 7.1. Let \mathcal{F} be the frontier of T in V . Lemma 5.4 implies that $D_{r_1}(P) = D_{r_1}(T, (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}}))$, where $\mathbf{d}_{\mathcal{F}}$ and $\mathbf{d}_{\mathcal{E}} \in Ad(\mathbf{d}_{\mathcal{F}})$ are defined

there. By the construction there and our hypothesis, $d_e \geq d_1$ for each e with $(e, v) \in \mathcal{F}$ for some v , so appealing to Table 2, we find that $D_{r_1}(P) > 1.15527$. This again contradicts the fact that $F - U$ has area less than 1.07.

By the above, no 2-cell of the centered dual decomposition has five edges. If a centered dual 2-cell Q with $n > 5$ edges is a centered Delaunay polygon, then by [3, Corollary 5.8] $D_{r_1}(Q) = D_{r_1}(P_n(d_1))$. This increases with n , by [3, Lemma 6.6], so an appeal to Table 1 establishes a contradiction as above. Now assume that Q contains a component T of $V_n^{(1)}$, let \mathcal{F} be the frontier of T in $V^{(1)}$, and define $\mathbf{d}_{\mathcal{F}}$ and $\mathbf{d}_{\mathcal{E}} \in Ad(\mathbf{d}_{\mathcal{F}})$ as in Lemma 5.4.

Again $d_e \geq d_1$ for each $e \in \mathcal{E}$ or such that $(e, v) \in \mathcal{F}$ for some v . Thus if the root vertex v_T of T has valence at least 5 in $V^{(1)}$ then:

$$D_{r_1}(T, \mathbf{d}) \geq D_{r_1}(P_{v_T}(\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})) \geq D_{r_1}(P_5(d_1)) > 1.22$$

If v_T has valence four in $V^{(1)}$ and a vertex v of T adjacent to v_T has valence at least four, let $T_0 = e_v \subset T$ have frontier \mathcal{F}_0 in $V^{(1)}$. Applying the basic algorithm with $\mathbf{b}_{\mathcal{F}_0} = (d_1, \dots, d_1)$ yields a bound of 1.8623. Proposition 6.3 thus implies that $D_{r_1}(T, \mathbf{d}) > 1.8623$.

If v_T has valence three in $V^{(1)}$, so does a vertex v adjacent to v_T in T , and a vertex $w \neq v_T$ adjacent to v in T has valence at least four, applying the basic algorithm to $T_0 = e_v \cup e_w$ with $\mathbf{b}_{\mathcal{F}_0} = (d_1, \dots, d_1)$ yields a bound of 2.46104, so Proposition 6.3 implies that $D_{r_1}(T, \mathbf{d}) > 2.46104$. In all other cases T has a subtree T_0 , containing v_T , with the same combinatorics and frontier in $V^{(1)}$ as a tree from Figure 7.1, so as above $D_{r_1}(T, \mathbf{d}) > 1.15$. In all cases we obtain a contradiction. \square

Lemma 7.3. *Let $T = e$ have frontier $\mathcal{F} = \{(e_0, v), (e_1, v), (e_2, v_T), (e_3, v_T)\}$, where v_T is the root vertex and v is the other. For d_β as in Example 2.4 and $\mathbf{d}_{\mathcal{F}} \geq (d_\beta, d_\beta, d_\beta, d_\beta)$, $D_{r_\beta}(T, (\mathbf{d}_{\mathcal{E}}, \mathbf{d}_{\mathcal{F}})) > D_{r_\beta}(P_4(d_\beta))$ for each $\mathbf{d}_{\mathcal{E}} \in Ad(\mathbf{d}_{\mathcal{F}})$, where $r_\beta = d_\beta/2$.*

Proof. We first note that if $Ad(\mathbf{d}_{\mathcal{F}}) \neq \emptyset$, where $\mathbf{d}_{\mathcal{F}} = (d_0, d_1, d_2, d_3)$, then for $\mathbf{d}_{\mathcal{E}} = d_e \in Ad(\mathbf{d}_{\mathcal{F}})$, $P_v(\mathbf{d}) = (d_e, d_0, d_1) \in \widehat{\mathcal{AC}}_3 - \widehat{\mathcal{C}}_3$ and $P_{v_T}(\mathbf{d}) = (d_e, d_2, d_3) \in \widehat{\mathcal{C}}_3$ by Definition 5.2. Thus [3, Corollary 3.11] and [3, Lemma 6.2] imply that:

$$\cosh d_0 + \cosh d_1 - 1 = b_0(d_0, d_1) \leq \cosh d_e < b_0(d_2, d_3) = \cosh d_2 + \cosh d_3 - 1$$

In particular, $\cosh d_0 + \cosh d_1 < \cosh d_2 + \cosh d_3$, so if $\mathbf{d}_{\mathcal{F}} \geq (d_\beta, d_\beta, d_\beta, d_\beta)$ then at least one of d_2 or d_3 is properly larger than d_β .

For $\mathbf{d}_{\mathcal{F}} \geq (d_\beta, d_\beta, d_\beta, d_\beta)$ with $Ad(\mathbf{d}_{\mathcal{F}}) \neq \emptyset$, let us first suppose that the minimum of $D_{r_\beta}(T, (d_e, \mathbf{d}_{\mathcal{F}}))$ over $Ad(\mathbf{d}_{\mathcal{F}})$ occurs at $(d^-, \mathbf{d}_{\mathcal{F}})$ satisfying Case (2) of Lemma 5.9. The corresponding improvement of the basic algorithm, found in Section 6.2, outputs 0.74844 given $R = r_2$ satisfying $\cosh r_2 = 2.8299$ and $\mathbf{b}_{\mathcal{F}} = (d_1, d_1, d_1, d_1)$. Since $r_2 > r_\beta$ and $d_1 < d_\beta$, it follows that $D_{r_\beta}(d_e, \mathbf{d}_{\mathcal{F}}) > 0.74844$ for each $d_e \in Ad(\mathbf{d}_{\mathcal{F}})$. On the other hand, $D_{r_\beta}(P_4(d_\beta)) < D_{r_1}(P_4(d_2)) < 0.56596$, where $\cosh d_2 = 15.0167$, so $D_{r_\beta}(d_e, \mathbf{d}_{\mathcal{F}}) > D_{r_\beta}(P_4(d_\beta))$ in this case.

Now suppose that the minimum of $D_{r_\beta}(T, (d_e, \mathbf{d}_{\mathcal{F}}))$ over $Ad(\mathbf{d}_{\mathcal{F}})$ occurs at $(d^-, \mathbf{d}_{\mathcal{F}})$ satisfying Case (1) of Lemma 5.9. Then $d^- = b_0(d_0, d_1) < b_0(d_2, d_3)$ by the first paragraph above, so $M_{r_\beta}(v_T, \mathbf{d}_{\mathcal{F}})$ as defined in Lemma 6.4 is equal to $D_{r_\beta}(d^-, d_2, d_3)$. Given $R = r_\beta$ and

$\mathbf{d}_{\mathcal{F}}$, the corresponding improvement of the basic algorithm (from Section 6.1) thus outputs $D_{r_{\beta}}(d^-, d_0, d_1) + D_{r_{\beta}}(d^-, d_2, d_3)$.

[3, Lemma 6.8] implies that $D_{r_{\beta}}(P_4(d_{\beta})) = 2D_{r_{\beta}}(B_0, d_{\beta}, d_{\beta})$, where $B_0 = b_0(d_{\beta}, d_{\beta})$. Since $\mathbf{d}_{\mathcal{F}} \geq (d_{\beta}, d_{\beta}, d_{\beta}, d_{\beta})$, the monotonicity property of b_0 recorded in [3, Lemma 3.4] implies that $d^- \geq B_0$. Therefore $(d^-, d_0, d_1) \geq (B_0, d_{\beta}, d_{\beta})$ and $(d^-, d_2, d_3) \geq (B_0, d_{\beta}, d_{\beta})$ in the sense of [3, Definition 5.7], and by the first paragraph above the latter inequality is proper. Thus [3, Corollary 5.8] implies the lemma. \square

Theorem 0.1. *Let $r_{\beta} = d_{\beta}/2 > 0$, where $\cosh d_{\beta}$ is the real root of $x^3 - 14x^2 - 15x - 4$. The Delaunay tessellation of a closed, orientable hyperbolic surface F of genus 2 determined by $\{x\}$ has all edges centered if F has injectivity radius $r \geq r_{\beta}$ at x . It is a triangulation unless $r = r_{\beta}$ and each edge has length d_{β} , in which case it has a single quadrilateral 2-cell.*

Proof. Suppose that a genus-two surface F has injectivity radius at least r_{β} at a point x . Since $2r_{\beta} = d_{\beta} > d_1$, Corollary 7.2 implies that no 2-cell of the centered dual tessellation determined by $\{x\}$ has more than four edges.

A hyperbolic disk U with radius r_{β} is embedded in F centered at x , and since $r_{\beta} > r_1$, the complementary area to U in F is less than 1.07 (see the first paragraph of the proof of Corollary 7.2). Each edge of the centered dual decomposition has length at least d_{β} , so if P is a centered quadrilateral 2-cell of this decomposition then $P \geq P_4(d_{\beta})$ in the sense of [3, Definition 5.7], and by [3, Corollary 5.8] $D_{r_{\beta}}(P) \geq D_{r_{\beta}}(P_4(d_{\beta}))$.

For a quadrilateral 2-cell P containing a component T of $V_n^{(1)}$, where V is the Voronoi tessellation determined by $\{x\}$, an argument like that for Remark 7.1 implies that T and its frontier \mathcal{F} are as described in Lemma 7.3. The conclusion there and Lemma 5.4 thus imply that $D_{r_{\beta}}(P) > D_{r_{\beta}}(P_4(d_{\beta}))$ in this case.

If r_2 satisfies $\cosh r_2 = 2.8299$ then $r_2 > r_{\beta}$, so since $d_{\beta} > d_1$ we have $D_{r_{\beta}}(P_4(d_{\beta})) > D_{r_2}(P_4(d_1)) > 0.56573$. Thus by the above the centered dual tessellation has at most one quadrilateral 2-cell, since $0.56573 \cdot 2 = 1.13146 > 1.07$ and by Proposition 4.2, $D_{r_{\beta}}(P)$ is the area of $P - (P \cap U)$ for each 2-cell P .

An Euler characteristic calculation now reveals that the centered dual tessellation of F determined by $\{x\}$ consists of either six triangles or a quadrilateral and four triangles. Recall that the tessellated surface F_{β} of Example 2.4 has the latter combinatorics and all edges of length d_{β} . By Corollary 2.11 this is the Delaunay tessellation of F_{β} determined by $\{x_{\beta}\}$, and since each polygon is centered by construction, it is also the centered dual tessellation. Since F_{β} has area 4π , Proposition 4.2 thus gives:

$$D_{r_{\beta}}(P_4(d_{\beta})) + 4 \cdot D_{r_{\beta}}(P_3(d_{\beta})) = 4\pi - 2\pi(\cosh r_{\beta} - 1)$$

If the centered dual tessellation of F determined by $\{x\}$ has a quadrilateral 2-cell P , let T_1, T_2, T_3 , and T_4 be its triangular 2-cells. Remark 7.1 implies that no triangular 2-cell of the centered dual decomposition contains a component of $V_n^{(1)}$; hence each is a centered Delaunay polygon. Therefore since each edge of the centered dual tessellation has length at least d_{β} ,

$D_{r_\beta}(T_i) \geq D_{r_\beta}(P_3(d_\beta))$ for each i by [3, Corollary 5.8]. By the above $D_{r_\beta}(P) \geq D_{r_\beta}(P_4(d_\beta))$, with strict inequality if P contains a component of $V_n^{(1)}$. In the latter case:

$$(7.3.1) \quad D_{r_\beta}(P) + \sum_{i=1}^4 D_{r_\beta}(T_i) > 4\pi - 2\pi \cdot (\cosh r_\beta - 1)$$

But this contradicts Proposition 4.2 and the fact that F has area 4π .

It follows that each 2-cell of the centered dual decomposition determined by $\{x\}$ is centered, and hence that this is also the Delaunay tessellation. If F has injectivity radius greater than r_β at x then each edge of the Delaunay tessellation has length greater than d_β . Thus if the Delaunay tessellation had quadrilateral component P in this case, we would again have the inequality (7.3.1). This is again a contradiction, and the theorem follows. \square

8. GEOMETRIC CONSEQUENCES

This section describes the geometric consequences of Theorem 0.1 for hyperbolic surfaces of genus 2 that have large injectivity radius at some point. It will be convenient to work with a model space that we may regard as parametrizing the set of pairs (F, x) , for F in \mathcal{M}_2 and $x \in F$, using the Delaunay tessellation determined by $\{x\}$.

Definition 8.1. We will say that an *edge-pairing* is a fixed point-free involution $\iota \in S_{18}$. For an edge-pairing ι , define

$$\mathcal{P}_\iota = \left\{ (d_0, \dots, d_{17}) \in (\widetilde{\mathcal{AC}}_3)^6 \mid d_i = d_{\iota(i)} \ \forall i, \sum_{j=0}^5 D_0(d_{3j}, d_{3j+1}, d_{3j+2}) = 4\pi \right\}$$

For $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota$ and $j \in \{0, \dots, 5\}$, let $T_j \subset \mathbb{H}^2$ be represented by $(d_{3j}, d_{3j+1}, d_{3j+2})$ in the sense of [3, Definition 3.1], with sides γ_i such that $\ell(\gamma_i) = d_i$ for $i \in \{3j, 3j+1, 3j+2\}$. Let $F_\iota(d_0, \dots, d_{17})$ be obtained from $\bigsqcup_{j=0}^5 T_j$ by isometrically identifying γ_i with $\gamma_{\iota(i)}$ for each i so that $x_{i-1} \rightarrow x_{\iota(i)}$ and $x_i \rightarrow x_{\iota(i)-1}$ for each $i \in \{3j, 3j+1, 3j+2\}$, and induce a metric on $F_\iota(d_0, \dots, d_{17})$ from those on the T_j (in the sense of, say, [2, Ch. I.7]).

Let us make a few initial observations:

- By the definition of d_α in Example 2.2, $(d_\alpha, \dots, d_\alpha) \in \mathcal{P}_\iota$ for each edge-pairing ι .
- The cellular isomorphism type of $F_\iota(d_0, \dots, d_{17})$ does not depend on the choice of $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota$. We say ι is *one-vertex* if (say) $F_\iota(d_\alpha, \dots, d_\alpha)$ has one vertex.
- For each one-vertex edge-pairing ι and $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota$, $F_\iota(d_0, \dots, d_{17})$ is isometric to a closed, orientable hyperbolic surface of genus 2.

The final observation above follows from the fact that the triangles T_j have total angle sum 2π , so the single vertex of $F_\iota(d_0, \dots, d_{17})$ has a neighborhood isometric to one in \mathbb{H}^2 . We may thus take $(d_0, \dots, d_{17}) \mapsto F_\iota(d_0, \dots, d_{17})$ as defining a map $F_\iota: \mathcal{P}_\iota \rightarrow \mathcal{M}_2$. To show continuity of this map we will lift it to \mathcal{T}_2 .

Definition 8.2. For a one-vertex edge-pairing ι , fix a maximal subtree S of the one-skeleton of the abstract dual to $F_\iota = F_\iota(d_\alpha, \dots, d_\alpha)$, and let $[\gamma_{i_0}], [\gamma_{i_1}], [\gamma_{i_2}], [\gamma_{i_3}]$ be the edges of $F_\iota^{(1)}$ not dual to edges of S . Further fix $v \in \mathbb{H}^2$ and a geodesic ray δ from v .

For $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota$ embed T_0 in \mathbb{H}^2 with its center at v and $x_0 = \gamma_0 \cap \gamma_1 \in \delta$, and for $j > 0$ embed T_j in \mathbb{H}^2 so that $T_j \cap T_{j'} = \gamma_i$ for each $i \in \{3j, 3j+1, 3j+2\}$ such that $[\gamma_i] \subset F_\iota(d_0, \dots, d_{17})$ is dual to an edge of S , where $\iota(i) \in \{3j', 3j'+1, 3j'+2\}$. Let $O = \bigcup_{j=0}^5 T_j \subset \mathbb{H}^2$. For $j \in \{0, 1, 2, 3\}$ let $f_j \in \text{Isom}^+(\mathbb{H}^2)$ satisfy $f(\gamma_{\iota(i_j)}) = \gamma_{i_j}$ and $f(x_{\iota(i_j)}) = x_{i_j-1}$, and let $\tilde{F}_\iota(d_0, \dots, d_{17}) = (f_0, f_1, f_2, f_3) \subset (\text{Isom}^+(\mathbb{H}^2))^4$.

That an embedding of T_0 is prescribed by the choice of v and δ , given $(d_0, d_1, d_2) \in \widetilde{\mathcal{AC}}_3$, follows from [3, Proposition 4.10]. Since S is a tree the embedding of T_0 and the criteria of Definition 8.2 determine embeddings of the other T_j . Then $O = \bigcup_{j=0}^5 T_j$ is an octagon with edge set $\{\gamma_{i_j}, \gamma_{\iota(i_j)}\}_{j=0}^3$, and the Poincaré polygon theorem implies that $\langle f_j \rangle_{j=0}^3$ is a discrete subgroup of $\text{Isom}^+(\mathbb{H}^2)$ with fundamental domain O and quotient isometric to $F_\iota(d_0, \dots, d_{17})$.

Fixing generators g_0, g_1, g_2, g_3 for $\pi_1 F$, an embedding $\text{Hom}(\pi_1 F, \text{Isom}^+(\mathbb{H}^2)) \hookrightarrow (\text{Isom}^+(\mathbb{H}^2))^4$ is given by $\rho \mapsto (\rho(g_0), \rho(g_1), \rho(g_2), \rho(g_3))$. \mathcal{T}_2 inherits the *algebraic topology* as the subspace topology from this embedding, when it is identified with the set of discrete, faithful representations in $\text{Hom}(\pi_1 F, \text{Isom}^+(\mathbb{H}^2))$ (see [4, §10.3]). The paragraph above implies that $\tilde{F}_\iota(d_0, \dots, d_{17}) \in \tilde{\mathcal{T}}_2$. Therefore $(d_0, \dots, d_{17}) \mapsto \tilde{F}_\iota(d_0, \dots, d_{17})$ determines a map $\tilde{F}_\iota: \mathcal{P}_\iota \rightarrow \mathcal{T}_2$. Since $F_\iota(d_0, \dots, d_{17})$ is the quotient of \mathbb{H}^2 by $\langle f_j \rangle_{j=0}^3$, \tilde{F}_ι lifts F_ι .

Lemma 8.3. *For each one-vertex edge-pairing ι , $\tilde{F}_\iota: \mathcal{P}_\iota \rightarrow \mathcal{T}_2$ is continuous, where \mathcal{T}_2 has the algebraic topology.*

Proof. It is well known that an orientation-preserving isometry of \mathbb{H}^2 is determined by its values on distinct $x, y \in \mathbb{H}^2$, and that the isometry so-determined varies continuously with the destinations of x and y . The following claim thus implies the lemma: the vertices of the octagon O from Definition 8.2 vary continuously with $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota$.

We will show that the vertices of each T_j , embedded as prescribed in Definition 8.2, vary continuously with (d_0, \dots, d_{17}) . We use induction, outward on the tree S from its vertex corresponding to T_0 . The base case follows directly from [3, Proposition 4.10], which implies that the vertices of T_0 vary continuously with (d_0, d_1, d_2) .

At least one edge of T_0 , say γ_0 , is dual to an edge of S . Then for j such that T_j contains $\gamma_{\iota(0)}$, we embed T_j in \mathbb{H}^2 as prescribed in Definition 8.2 in two steps: first embed T_j with center v , $x_{\iota(0)} \in \delta$, and $\gamma_{\iota(0)} \subset \mathcal{H}$, then move it via an isometry so that $\gamma_{\iota(0)} = \gamma_0$ and $x_{\iota(0)} = x_1$. By [3, Proposition 4.10], the vertices of the initial embedding vary continuously with (d_0, \dots, d_{17}) , so by the base case and the observation at the beginning of this proof, the vertices of the second do as well. The general inductive step is no more complicated. \square

Since the moduli space \mathcal{M}_2 inherits its usual topology as the quotient of a discontinuous group action on \mathcal{T}_2 (again see [4, §10.3]), Lemma 8.3 implies that $F_\iota: \mathcal{P}_\iota \rightarrow \mathcal{M}_2$ is continuous for each one-vertex edge-pairing ι .

We now fix attention on $\mathcal{M}_2^{(r_\beta)}$, the set of $F \in \mathcal{M}_2$ with injectivity radius at least r_β at some $x \in F$, where r_β is as defined in Example 2.4. Such pairs (F, x) fall under the purview of Theorem 0.1. We will begin by relating $\mathcal{M}_2^{(r_\beta)}$ to the model spaces \mathcal{P}_ι .

Lemma 8.4. *For $r \in [r_\beta, r_\alpha]$, where r_α is as in Example 2.2, let $d_r = 2r$. There is a unique $d_1(r) \in [d_r, b_0(d_r, d_r)]$, where b_0 is as in [3, Lemma 3.4], so that:*

$$4\pi = 4 \cdot D_{0,3}(d_1(r)) + 2 \cdot D_0(d_1(r), d_r, d_r)$$

This satisfies $d_1(r_\beta) = b_0(d_\beta, d_\beta)$, $d_1(r_\alpha) = d_\alpha$, and $d_1(r) \in (d_r, b_0(d_r, d_r))$ for $r \in (r_\beta, r_\alpha)$. For a one-vertex edge-pairing ι and $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota \cap (\tilde{\mathcal{C}}_3 \cup \widetilde{\mathcal{BC}}_3)^6$ such that $d_i \geq d_r$ for each i , $d_i \leq d_1(r)$ for each i .

Proof. For each $d \in [d_r, b_0(d_r, d_r)]$, [3, Corollary 3.11] implies that $(d, d_r, d_r) \in \tilde{\mathcal{C}}_3 \cup \widetilde{\mathcal{BC}}_3$. [3, Proposition 5.5] thus implies that $D_0(d, d_r, d_r)$ increases with d on this interval, so:

$$6 \cdot D_{0,3}(d_r) \leq 4 \cdot D_{0,3}(d_r) + 2 \cdot D_0(d, d_r, d_r) \leq 4 \cdot D_{0,3}(d_r) + 2 \cdot D_0(b_0(d_r, d_r), d_r, d_r)$$

Since $d_r \leq d_\alpha$, [3, Proposition 5.5] and the construction of d_α (see Example 2.2) imply that $6 \cdot D_{0,3}(d_r) \leq 4\pi$, with equality if and only if $d_r = d_\alpha$. As we observed in Example 2.12, for d_β as defined in Example 2.4, [3, Lemma 6.8] implies that $4 \cdot D_{0,3}(d_\beta) + 2 \cdot D_0(b_\beta, d_\beta, d_\beta) = 4\pi$, where $b_\beta = b_0(d_\beta, d_\beta)$. Therefore since $d_\beta \leq d_r$, [3, Proposition 5.5] gives:

$$4\pi = 4 \cdot D_{0,3}(d_\beta) + 2 \cdot D_0(b_\beta, d_\beta, d_\beta) \leq 4 \cdot D_{0,3}(d_r) + 2 \cdot D_0(b_0(d_r, d_r), d_r, d_r),$$

with equality if and only if $d_r = d_\beta$. The continuity and monotonicity of D_0 imply that $d_1(r)$ exists and is unique for each r . Furthermore, by the above $d_r < d_1(r) < b_0(d_r, d_r)$ unless $r = r_\beta$ or $r = r_\alpha$, and $d_1(r_\beta) = b_\beta$ and $d_1(r_\alpha) = d_\alpha$.

For a one-vertex edge-pairing ι and $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota \cap (\tilde{\mathcal{C}}_3 \cup \widetilde{\mathcal{BC}}_3)^6$ such that $d_i \geq d_r$ for each i , [3, Corollary 5.8] implies that $D_0(d_{3j}, d_{3j+1}, d_{3j+2}) \geq D_{0,3}(d_r)$ for each j between 0 and 5, where $D_{0,3}$ is as in [3, Lemma 6.6]. If $d = \max\{d_i\}_{i=0}^{17}$ then another application of [3, Corollary 5.8] gives:

$$\sum_{i=1}^6 D_0(T_i) \geq 4 \cdot D_{0,3}(d_r) + 2 \cdot D_0(d, d_r, d_r)$$

To justify the “2” above, note that if $d_i = d$ then $d_{\iota(i)} = d$ as well, and there is no j such that $i, \iota(i) \in \{3j, 3j+1, 3j+2\}$ since ι is one-vertex. If $d > d_1(r)$, then by construction (and [3, Corollary 5.8]) it would follow that $\sum_{i=1}^6 D_0(T_i) > 4\pi$, contradicting $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota$. \square

Lemma 8.5. *If $F \in \mathcal{M}_2$ has injectivity radius $r \geq r_\beta$ at x , then there is a one-vertex edge-pairing ι and $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota \cap (\tilde{\mathcal{C}}_3 \cup \widetilde{\mathcal{BC}}_3)^6 \cap [d_\beta, b_0(d_\alpha, d_\alpha)]^{18}$, such that F is isometric to $F_\iota(d_0, \dots, d_{17})$, taking x to its vertex.*

Proof. For such $F \in \mathcal{M}_2^{(r_\beta)}$ and $x \in F$, let P be the Delaunay tessellation of F determined by $\{x\}$. By Theorem 0.1 and Lemma 3.3, each 2-cell of P is centered. If P is a triangulation then it has 2-cells T_0, \dots, T_5 . Cyclically ordering the sides of each T_j and recording their lengths,

determines a one-vertex edge-pairing ι and $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota$ with $F = F_\iota(d_0, \dots, d_{17})$ and $(d_{3j}, d_{3j+1}, d_{3j+2}) \in \widetilde{\mathcal{C}}_3$ for each j .

If P has a quadrilateral 2-cell O , a diagonal of O divides it into two cyclic triangles. Then making choices as above yields ι and (d_0, \dots, d_{17}) with $F = F_\iota(d_0, \dots, d_{17})$. In this case, however, for each j corresponding to a triangle in O the corresponding triple $(d_{3j}, d_{3j+1}, d_{3j+2})$ is in $\widetilde{\mathcal{BC}}_3$.

For each i , $d_i \geq d_\beta = 2r_\beta$ since it is an edge length of P , so Lemma 8.4 implies that $d_i \leq d_1(r) \leq b_0(d_r, d_r)$ for each i . Furthermore, $b_0(d_r, d_r) \leq b_0(d_\alpha, d_\alpha)$ by [3, Lemma 3.4], since $d_r \leq d_\alpha$ by Lemma 2.3. \square

Lemma 8.6. *For each one-vertex edge-pairing ι and $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota \cap (\widetilde{\mathcal{C}} \cup \widetilde{\mathcal{BC}}_3)^6$, $F \doteq F_\iota(d_0, \dots, d_{17})$ has injectivity radius $r = \min\{d_i/2\}_{i=0}^{17}$ at its vertex x .*

Proof. We argue as in Examples 2.2 and 2.4, using [3, Lemma 5.3]. Of course F has injectivity radius at most r at x . Since each $(d_{3j}, d_{3j+1}, d_{3j+2}) \in \widetilde{\mathcal{C}}_3 \cup \widetilde{\mathcal{BC}}_3$ for each j , each T_j contains the entire sector that it determines in a disk with radius r centered at any of its vertices. Since the T_j have total angle measure 2π , a disk with radius r is embedded in F , centered at x . \square

Corollary 8.7. $M_2^{(r_\beta)}$ is compact.

Proof. Lemmas 8.5 and 8.6 together imply that $\mathcal{M}_2^{(r_\beta)}$ is the union, taken over the finite collection of one-vertex edge-pairings ι , of the image of $\mathcal{P}_\iota \cap (\widetilde{\mathcal{C}}_3 \cup \widetilde{\mathcal{BC}}_3)^6 \cap [d_\beta, b_0(d_\alpha, d_\alpha)]^{18}$ under F_ι . Each such F_ι is continuous by Lemma 8.3. $\widetilde{\mathcal{C}}_3 \cup \widetilde{\mathcal{BC}}_3$ is closed in \mathbb{R}^+ by [3, Lemma 3.3] and [3, Lemma 3.4], so since $d_\beta > 0$ its intersection with $[d_\beta, b_0(d_\alpha, d_\alpha)]^3$ is compact. The Corollary follows. \square

Definition 8.8. Say the *covering radius* of $\mathcal{S} \subset \mathbb{H}^2$ is $\inf\{r \leq \infty \mid \mathbb{H}^2 \subset \bigcup_{x \in \mathcal{S}} B_r(x)\}$.

Lemma 8.9. *Let F be a closed surface with universal cover $p: \mathbb{H}^2 \rightarrow F$, fix $\mathcal{S} \subset F$ finite, and let V be the Voronoi tessellation of F determined by \mathcal{S} . Then $\widetilde{\mathcal{S}} \doteq p^{-1}(\mathcal{S})$ has covering radius equal to $\max\{J_v \mid v \in V^{(0)}\}$, where J_v is as in Lemma 1.2.*

Proof. It is clear that the covering radius of $\widetilde{\mathcal{S}}$ is at least the quantity above, since for any $v \in \widetilde{V}^{(0)}$, $d(x, v) \geq J_{p(v)}$ for each $x \in \widetilde{\mathcal{S}}$, where $\widetilde{V} = p^{-1}(V)$. We will prove equality by showing that $P_v \subset \bigcup_{i=0}^{n-1} \overline{B}_{J_{p(v)}}(x_i)$ for each such v , where P_v as in Lemma 1.9 has vertex set $\{x_i\}_{i=0}^{n-1}$, and $\overline{B}_r(x)$ is the closed ball of radius r about x .

Let us assume that the $\{x_i\}$ are cyclically ordered in the sense of Definition 1.5, and for each i let γ_i be the side of P_v bounded by x_{i-1} and x_i (with $i-1$ taken modulo n). Since each endpoint of γ_i lies on a circle with radius $J_{p(v)}$, $\ell(\gamma_i) \leq 2J_{p(v)}$. Therefore the midpoint m_i of γ_i is in $\overline{B}_{J_{p(v)}}(x_{i-1}) \cap \overline{B}_{J_{p(v)}}(x_i)$.

Since the convex set $\overline{B}_{J_{p(v)}}(x_i)$ contains x_i , m_i , and v , it contains the right triangle in \mathbb{H}^2 that they determine. Similarly, $B_{J_{p(v)}}(x_{i-1})$ contains the triangle determined by x_{i-1} , m_i , and v .

The union of these triangles is T_i as defined in [3, Lemma 1.6], so $T_i \subset \overline{B}_{J_{p(v)}}(x_{i-1}) \cup \overline{B}_{J_{p(v)}}(x_i)$. Since this holds for each i , [3, Lemma 1.6] implies that $P_v \subset \bigcup_{i=0}^{n-1} \overline{B}_{J_{p(v)}}(x_i)$. \square

For a surface F with universal cover $p: \mathbb{H}^2 \rightarrow F$ and $x \in F$, it is clear that the covering radius of F at x (as defined below Theorem 0.2) is equal to the covering radius of $\tilde{\mathcal{S}} \doteq p^{-1}(x)$.

Definition 8.10. For a one-vertex edge-pairing ι , say $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota$ is *exceptional* if $d_i = d_\beta$ for all but two $i \in \{0, \dots, 17\}$. (In this case $d_{i_0} = d_{\iota(i_0)} = b_0(d_\beta, d_\beta)$ for some i_0 .)

Lemma 8.11. For a one-vertex edge-pairing ι , if $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota$ is exceptional there exists i_0 such that $d_{i_0} = d_{\iota(i_0)} = b_0(d_\beta, d_\beta)$. For j_0 with $i_0 \in \{3j_0, 3j_0 + 1, 3j_0 + 2\}$ and j_1 with $\iota(i_0) \in \{3j_1, 3j_1 + 1, 3j_1 + 2\}$, $T_{j_0} \cup T_{j_1}$ is a 2-cell of the Delaunay tessellation P of $F_\iota(d_0, \dots, d_{17})$ determined by $\{x\}$, where x is the vertex. For $j \neq j_0, j_1$, T_j is a 2-cell of P .

Proof. [3, Proposition 5.5] implies that $x \mapsto D_0(x, d_\beta, d_\beta)$ achieves a unique maximum at $x = b_0(d_\beta, d_\beta)$. Since $4 \cdot D_{0,3}(d_\beta) + 2 \cdot D_0(b_0(d_\beta, d_\beta), d_\beta, d_\beta) = 4\pi$ by construction (see Example 2.12), it follows by definition of \mathcal{P}_ι that if $(d_0, \dots, d_{n-1}) \in \mathcal{P}_\iota$ is exceptional it must have two entries equal to $b_0(d_\beta, d_\beta)$. That they are exchanged by ι also follows by definition.

Note that j_0 and j_1 defined above are distinct, since ι is one-vertex. [3, Lemma 3.9] implies that T_{j_0} has its center in γ_{i_0} , and T_{j_1} has its in $\gamma_{\iota(i_0)}$, so their union is a centered quadrilateral with all edges of length d_β . Corollary 2.9 thus implies that $T_{j_0} \cup T_{j_1}$ is a 2-cell of P . It implies the same for T_j , $j \neq j_0, j_1$. \square

Lemma 8.12. For each one-vertex edge-pairing ι and non-exceptional $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota \cap (\tilde{\mathcal{C}}_3 \cup \widehat{\mathcal{BC}}_3)^6 \cap [d_\beta, b_0(d_\alpha, d_\alpha)]^{18}$, each T_j from Definition 8.1 is a 2-cell of the Delaunay tessellation of $F_\iota(d_0, \dots, d_{17})$ determined by $\{x\}$ where x is the vertex.

Proof. By Lemma 8.6, $F \doteq F_\iota(d_0, \dots, d_{17})$ has injectivity radius $r = \min\{d_i/2\}_{i=0}^{17}$ at its vertex x . By hypothesis $d_r = 2r = \min\{d_i\}_{i=0}^{17}$ is at least d_β , so Lemma 8.4 implies that $d_i \leq d_1(r) \leq b_0(d_r, d_r)$ for each i . If $d_{i_0} = b_0(d_r, d_r)$ for some i_0 , we claim that (d_0, \dots, d_{17}) is exceptional.

If $d_{i_0} = b_0(d_r, d_r)$ for some i_0 then in particular $d_1(r) = b_0(d_r, d_r)$, so $r = r_\beta$ by Lemma 8.4. Fix j_0 such that $i_0 \in \{3j_0, 3j_0 + 1, 3j_0 + 2\}$. Since ι is one-vertex, $j_0 \neq j'_0$ such that $\iota(i_0) \in \{3j'_0, 3j'_0 + 1, 3j'_0 + 2\}$. Applying [3, Corollary 5.8], we have:

$$\sum_{j=0}^5 D_0(3j, 3j+1, 3j+2) \geq 4 \cdot D_{0,3}(d_\beta) + D_0(b_0(d_\beta, d_\beta), d_\beta, d_\beta) = 4\pi$$

The latter equality is by construction (see Example 2.4). If there were $i \notin \{i_0, \iota(i_0)\}$ with $d_i > d_\beta$ then by [3, Corollary 5.8] the above inequality would be strict, contradicting $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota$. The claim follows.

For non-exceptional $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota$ the claim implies that $d_i < b_0(d_r, d_r)$ for each i . Thus $(d_{3j}, d_{3j+1}, d_{3j+2}) \in \tilde{\mathcal{C}}_3$ for each j , since for instance $d_{3j} < b_0(d_r, d_r) \leq b_0(d_{3j+1}, d_{3j+2})$ (the latter inequality follows from [3, Lemma 3.4]). Since $B_0 = b_0(d_r, d_r)$ satisfies $\cosh B_0 =$

$2 \cosh d_r - 1$ by [3, Lemma 6.2], Proposition 2.8 implies for each j that T_j is a 2-cell of the Delaunay tessellation determined by $\{x\}$. \square

Corollary 8.13. *For each one-vertex edge-pairing ι define $J_\iota: \mathcal{P}_\iota \rightarrow \mathbb{R}^+$ by $J_\iota(d_0, \dots, d_{17}) = \max\{J(d_{3j}, d_{3j+1}, d_{3j+2})\}_{j=0}^5$, where J is as in [3, Lemma 3.6]. For $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota \cap (\tilde{\mathcal{C}} \cup \widetilde{\mathcal{BC}}_3)^6 \cap [d_\beta, b_0(d_\alpha, d_\alpha)]^{18}$, $F_\iota(d_0, \dots, d_{17})$ has covering radius $J_\iota(d_0, \dots, d_{17})$ at its vertex x .*

Proof. For non-exceptional $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota \cap (\tilde{\mathcal{C}} \cup \widetilde{\mathcal{BC}}_3)^6 \cap [d_\beta, b_0(d_\alpha, d_\alpha)]^{18}$ this follows from Lemmas 8.9 and 8.12, using the bijective correspondence between the Voronoi tessellation's vertex set and the set of Delaunay 2-cells described in Lemma 1.9. Recall in particular that for each j , $J(d_{3j}, d_{3j+1}, d_{3j+2})$ is the radius of T_j by [3, Lemma 3.6]. If (d_0, \dots, d_{17}) is exceptional, Lemmas 8.9 and 8.11 combine in the same way to give the conclusion, noting additionally that $J(b_0(d_\beta, d_\beta), d_\beta, d_\beta) = J(d_\beta, d_\beta, d_\beta, d_\beta)$ by [3, Lemma 6.8]. \square

Lemma 8.14. *For a one-vertex edge-pairing ι and $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota \cap (\tilde{\mathcal{C}}_3 \cup \widetilde{\mathcal{BC}}_3)^6 \cap [d_\beta, b_0(d_\alpha, d_\alpha)]^{18}$, let $r = \min\{d_i/2\}$. If for each i $d_i < d_1(r)$, as defined in Lemma 8.4, then (d_0, \dots, d_{n-1}) deforms preserving $\min\{d_i\}_{i=0}^{17}$ but increasing $J_\iota(d_0, \dots, d_{17})$.*

Proof. Fix j_1 such that $J_\iota(d_0, \dots, d_{17}) = J(d_{3j_1}, d_{3j_1+1}, d_{3j_1+2})$ and $i_1 \in \{3j_1, 3j_1+1, 3j_1+2\}$ such that $d_{i_1} = \max\{d_i\}_{i=3j_1}^{3j_1+2}$. Let j'_1 be such that $\iota(i_1) \in \{3j'_1, 3j'_1+1, 3j'_1+2\}$. Since ι is one-vertex, $j_1 \neq j'_1$. There exists $i_0 \notin \{i_1, \iota(i_1)\}$ with $d_{i_0} > d_r = 2r$, since otherwise by [3, Corollary 5.8] and the definition of $d_1(r)$:

$$\sum_{j=0}^5 D_0(d_{3j}, d_{3j+1}, d_{3j+2}) = 4 \cdot D_{0,3}(d_r) + 2 \cdot D_0(d_{i_1}, d_r, d_r) < 4\pi$$

This would contradict $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota$. Fix j_0 with $i_0 \in \{3j_0, 3j_0+1, 3j_0+2\}$ and j'_0 with $\iota(i_0) \in \{3j'_0, 3j'_0+1, 3j'_0+2\}$, and note as above that $j_0 \neq j'_0$. We will deform (d_0, \dots, d_{n-1}) changing only $d_{i_0} = d_{\iota(i_0)}$ and $d_{i_1} = d_{\iota(i_1)}$.

Suppose first that $\{j_1, j'_1\} = \{j_0, j'_0\}$. (In this case T_{j_1} and $T_{j'_1}$ share edges corresponding to γ_{i_1} and γ_{i_0} in $F_\iota(d_0, \dots, d_{17})$.) Let d' be the element of $\{d_{3j_1}, d_{3j_1+1}, d_{3j_1+2}\}$ not equal to d_{i_1} or d_{i_0} , and let d'' be the corresponding element of $\{d_{3j'_1}, d_{3j'_1+1}, d_{3j'_1+2}\}$. For small $t \geq 0$, we will take $d_{i_1}(t) = d_{\iota(i_1)}(t) = d_{i_1} + t$ and choose $d_{i_0}(t) = d_{\iota(i_0)}(t)$ so that $D_0(d_{i_1} + t, d_{i_0}(t), d') + D_0(d_{i_1} + t, d_{i_0}(t), d'')$ is constant. By [3, Proposition 5.5], $d_{i_0}(t)$ must satisfy:

$$\frac{d}{dt}(d_{i_0}(t)) = -\frac{\sqrt{\frac{1}{\cosh^2((d_{i_1}+t)/2)} - \frac{1}{\cosh^2 J_1(t)}} + \sqrt{\frac{1}{\cosh^2((d_{i_1}+t)/2)} - \frac{1}{\cosh^2 J'_1(t)}}}{\sqrt{\frac{1}{\cosh^2(d_{i_0}(t)/2)} - \frac{1}{\cosh^2 J_1(t)}} + \sqrt{\frac{1}{\cosh^2(d_{i_0}(t)/2)} - \frac{1}{\cosh^2 J'_1(t)}}}$$

Above $J_1(t) = J(d_{i_1} + t, d_{i_0}(t), d')$ and $J'_1(t) = J(d_{i_1} + t, d_{i_0}(t), d'')$. The existence theorem for ordinary differential equations implies that a unique differentiable function $d_{i_0}(t)$, defined on $[0, \epsilon)$ for some $\epsilon > 0$, satisfies the equation above. Using this equation we find that $d_{i_0}(t)$ decreases in t , and also that $|d'_{i_0}(t)| < 1$, since it follows that $d_{i_0}(t) < d_{i_1} + t$ for all $t \in [0, \epsilon)$.

Since $d_{i_0}(t) < d_{i_1} + t$ for all $t > 0$, [3, Lemma 4.5] implies that $\left|\frac{\partial J}{\partial d_{i_0}(t)}\right| < \left|\frac{\partial J}{\partial(d_{i_1}+t)}\right|$, and since $|d'_{i_0}(t)| < 1$ the chain rule implies that $J_1(t)$, and hence also J_ι increases with t in this case.

There are three other possibilities for $\{j_1, j'_1, j_0, j'_0\}$: one in which all four of its elements are distinct and two in which it has only three distinct elements (we do not distinguish the case $j_1 = j_0$ from $j_1 = j'_0$, or $j'_1 = j_0$ from $j'_1 = j'_0$). In each case we change each element of this set by increasing d_{i_1} and decreasing d_{i_0} , leaving all other entries constant while keeping the defect sum unchanged.

As long as $i_0 \notin \{3j_1, 3j_1 + 1, 3j_1 + 2\}$ (equivalently, $j_1 \notin \{j_0, j'_0\}$) it is clear by [3, Lemma 4.5] that J_t increases with t , so it remains only to consider the case that $j_1 = j'_0$ but $j'_1 \neq j_0$. Taking $d_{i_1}(t) = d_{i_1} + t$, in this case $d_{i_0}(t)$ must satisfy the following differential equation:

$$\frac{d}{dt}(d_{i_0}(t)) = -\frac{\sqrt{\frac{1}{\cosh^2((d_{i_1}+t)/2)} - \frac{1}{\cosh^2 J_1(t)}} + \sqrt{\frac{1}{\cosh^2((d_{i_1}+t)/2)} - \frac{1}{\cosh^2 J'_1(t)}}}{\sqrt{\frac{1}{\cosh^2(d_{i_0}(t)/2)} - \frac{1}{\cosh^2 J_1(t)}} + \sqrt{\frac{1}{\cosh^2(d_{i_0}(t)/2)} - \frac{1}{\cosh^2 J_0(t)}}}$$

Here $J_1(t) = J(d_{i_0}(t), d_{i_1}+t, d')$, $J'_1(t) = J(d_{i_1}+t, d_{3j'_1+1}, d_{3j'_1+2})$ (assuming for simplicity that $\iota(i_1) = 3j'_1$), and $J_0(t) = J(d_{i_0}(t), d_{3j_0+1}, d_{3j_0+2})$ (assuming that $i_0 = 3j_0$). We may assume that all entries not in $\{d_{3j_1}, d_{3j_1+1}, d_{3j_1+2}\}$ equal d_r , since otherwise replacing d_{i_0} by an entry not in the set above allows appeal to another case. Thus in this case $J_0(t) = J(d_{i_0}(t), d_r, d_r)$ and $J'_1(t) = J(d_{i_1}+t, d_r, d_r)$.

Unlike the first case we considered, it is not immediately obvious here that $|d'_{i_0}(t)| < 1$: the problem is that the defect derivative function $\sqrt{\frac{1}{\cosh^2(d/2)} - \frac{1}{\cosh^2 J}}$ decreases in d but increases in J , and [3, Lemma 4.5] implies that $J_0(t) < J_1(t)$. However we have the following:

Claim 8.14.1. For fixed $d > 0$ and x such that $(x, d, d) \in \mathcal{C}_3$, the function

$$x \mapsto \frac{1}{\cosh^2(x/2)} - \frac{1}{\cosh^2 J(x, d, d)}$$

decreases in x .

Proof of claim. Simplifying the formula of [3, Lemma 6.1] gives:

$$\sinh J(x) = \frac{2 \sinh^2(d/2)}{\sqrt{4 \sinh^2(d/2) - \sinh^2(x/2)}}$$

Let us take $X = \cosh^2(x/2)$ and $D = \cosh^2(d/2)$. Inserting the formula above into the function in question, after some more simplification we obtain:

$$\frac{1}{X} - \frac{1}{\cosh^2 J(x, d, d)} = \frac{[2D - 1 - X]^2}{X[(2D - 1)^2 - X]} = \left[\frac{2D - 1}{X} - 1 \right] \left[1 - \frac{2(D - 1)(2D - 1)}{(2D - 1)^2 - X} \right]$$

By [3, Lemma 6.2], $(x, d, d) \in \mathcal{C}_3$ if and only if $X < 2D - 1$, thus for such x the functions in brackets on the right-hand side are positive-valued. Since they also clearly decrease with X , their product is decreasing. Since X increases with x , the claim is proved. \square

Since $d_{i_0}(t) < d_{i_1} + t$, the claim implies that $|d'_{i_0}(t)| < 1$, so as in the first case considered it follows that $J(d_{i_0}(t), d_{i_1} + t, d')$, and hence also J_t increases with t . In each case above we have thus produced deformations through \mathcal{P}_t so that no entries change but $d_{i_0} = d_{\iota(i_0)}$ and

$d_{i_1} = d_{\iota(i_1)}$ change. Furthermore, d_{i_1} increases with t and $d_{i_0}(0) = d_{i_0} > d_r$, so $d_0(t) > d_r = \min\{d_i\}_{i=0}^{17}$ for t near 0. \square

Lemma 8.15. *For each $r \in [r_\beta, r_\alpha]$ and one-vertex edge-pairing ι there exists $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota \cap (\widetilde{\mathcal{C}}_3 \cup \widetilde{\mathcal{BC}}_3)^6$ with $r = \min\{d_i/2\}_{i=0}^{17}$ and $d_{i_0} = d_1(r)$ for some i_0 , where $d_1(r)$ is as in Lemma 8.4. At each such point:*

$$\sinh J_\iota(d_0, \dots, d_{17}) = \frac{2 \sinh^2 r}{\sqrt{4 \sinh^2 r - \sinh^2(d_1(r)/2)}}$$

Proof. We may for instance take $d_0 = d_{\iota(0)} = d_1(r)$ and $d_i = d_r \doteq 2r$ for each $i \neq 0, \iota(0)$. Then $(d_0, \dots, d_{n-1}) \in \mathcal{P}_\iota \cap (\widetilde{\mathcal{C}}_3 \cup \widetilde{\mathcal{BC}}_3)^6$ by definition. The computation of J_ι above is a direct application of [3, Lemma 6.1], noting that Lemma 8.4 and [3, Corollary 5.8] imply that if $r = \min\{d_i/2\}_{i=0}^{17}$ and $d_{i_0} = d_1(r)$ for some i_0 then $d_i = d_r$ for all $i \neq i_0, \iota(i_0)$. \square

Proof of Theorem 0.2. For each $r \in [r_\beta, r_\alpha]$ and one-vertex edge-pairing ι , J_ι is a continuous function on \mathcal{P}_ι , being the maximum of functions which are themselves continuous by [3, Proposition 4.1]. Hence it attains a maximum on the following compact subset:

$$\bigcup_{j=0}^{17} \mathcal{P}_\iota \cap (\widetilde{\mathcal{C}}_3 \cup \widetilde{\mathcal{BC}}_3)^6 \cap ([d_r, b_0(d_r, d_r)]^j \times \{d_r\} \times [d_r, b_0(d_r, d_r)]^{17-j})$$

By Lemma 8.6, this consists of those $(d_0, \dots, d_{17}) \in \mathcal{P}_\iota \cap (\widetilde{\mathcal{C}}_3 \cup \widetilde{\mathcal{BC}}_3)^6$ such that $F_\iota(d_0, \dots, d_{17})$ has injectivity radius r at its vertex F . Lemma 8.14 implies that J_ι attains its maximum at such a point as described in Lemma 8.15, and the maximum is as described there. Since $d_1(r) \leq b_0(d_r, d_r)$ by Lemma 8.4, and $B_0 = b_0(d_r, d_r)$ satisfies $\sinh(B_0/2) = \sqrt{2} \sinh r$ by [3, Lemma 6.1], a simplification gives $\sinh J_\iota(d_0, \dots, d_{17}) \leq \sqrt{2} \sinh r$ for each (d_0, \dots, d_{17}) in the set above. The result now follows directly from Lemma 8.5 and Corollary 8.13. \square

REFERENCES

- [1] K. Böröczky. Packing of spheres in spaces of constant curvature. *Acta Math. Acad. Sci. Hungar.*, 32(3-4):243–261, 1978.
- [2] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [3] Jason DeBlois. The geometry of cyclic hyperbolic polygons. Preprint. arXiv:1101.4971v2, March 2011.
- [4] Benson Farb and Dan Margalit. *A primer on mapping class groups*. Princeton University Press, to appear.
- [5] David Mumford. A remark on Mahler’s compactness theorem. *Proc. Amer. Math. Soc.*, 28:289–294, 1971.
- [6] Jean-Marc Schlenker. Small deformations of polygons and polyhedra. *Trans. Amer. Math. Soc.*, 359(5):2155–2189, 2007.
- [7] Rolf Walter. Polygons in hyperbolic geometry 1: Rigidity and inversion of the n -inequality. Preprint. arXiv:1008.3404v1, August 2010.

DEPARTMENT OF MATHEMATICS, STANFORD UNIVERSITY

E-mail address: jdeblois@math.stanford.edu